Mon. Not. R. Astron. Soc. **000**, 000–000 (0000) Printed 20 December 2016 (MN LATEX style file v2.2)

Lyα emission line reconstruction for high-z QSOs

Bradley Greig[1⋆], Andrei Mesinger[1], Ian D. McGreer[2], Simona Gallerani[1], Zoltán Haiman[3]

[1]*Scuola Normale Superiore, Piazza dei Cavalieri 7, I-56126 Pisa, Italy*
[2]*Steward Observatory, The University of Arizona, 933 North Cherry Avenue, Tucson, AZ 85721-0065, USA*
[3]*Department of Astronomy, Columbia University, 550 West 120th Street, New York, NY 10027, USA*

arXiv:1605.09388v2 [astro-ph.GA] 19 Dec 2016

20 December 2016

ABSTRACT

We introduce an intrinsic Lyα emission line profile reconstruction method for high-z quasars (QSOs). This approach utilises a covariance matrix of emission line properties obtained from a large, moderate-z ($2 \leq z \leq 2.5$), high signal to noise (S/N$>$ 15) sample of BOSS QSOs. For each QSO, we complete a Monte Carlo Markov Chain fitting of the continuum and emission line properties and perform a visual quality assessment to construct a large database of robustly fit spectra. With this dataset, we construct a covariance matrix to describe the correlations between the high ionisation emission lines Lyα, C IV, Si IV +O IV] and C III], and find it to be well approximated by an N-dimensional Gaussian distribution. This covariance matrix characterises the correlations between the line width, peak height and velocity offset from systemic while also allowing for the existence of broad and narrow line components for Lyα and C IV. We illustrate how this covariance matrix allows us to statistically characterise the intrinsic Lyα line solely from the observed spectrum redward of 1275Å. This procedure can be used to reconstruct the intrinsic Lyα line emission profile in cases where Lyα may otherwise be obscured. Applying this reconstruction method to our sample of QSOs, we recovered the Lyα line flux to within 15 per cent of the measured flux at 1205Å (1220Å) $\sim$ 85 (90) per cent of the time.

Key words: quasars: emission lines – quasars: absorption lines – quasars: general – cosmology: observations – cosmology: theory

1 INTRODUCTION

The ability to measure the intrinsic spectrum of quasars (QSOs) plays an important role in astrophysics. Resolving individual emission line profiles can enable the detailed study of both the central supermassive black hole (SMBH) and the internal structure within the SMBH powered accretion disk (through reverberation mapping e.g. Blandford & McKee 1982; Peterson 1993 or the virial method e.g. Kaspi et al. 2000; Peterson et al. 2004; Vestergaard & Peterson 2006). On cosmological scales, the relative strength (or lack thereof) of the emission profiles and spectral slopes of the QSO continuum can be used to yield insights into the thermal and ionisation state of the intergalactic medium (IGM) through studies of the QSO proximity zone (e.g. Mesinger & Haiman 2004; Wyithe & Loeb 2004; Fan et al. 2006; Bolton & Haehnelt 2007a; Mesinger & Haiman 2007; Carilli et al. 2010; Calverley et al. 2011; Wyithe & Bolton 2011; Schroeder et al. 2013).

The key ingredient for using QSOs to explore the IGM is observing the intrinsic ultraviolet (UV) emission. Thermal emission from the accretion disk peaks in UV, which then interacts with surrounding neutral hydrogen and recombines two thirds of the time into Lyα photons (rest frame $\lambda_{\rm Ly\alpha} = 1215.67$Å) to produce a prominent Lyα emission profile. However, owing to the large scattering cross-section of Lyα photons, neutral hydrogen column densities of $N_{\rm H\,I} > 10^{18}{\rm cm}^{-2}$ are sufficient for Lyα to enter into the strong absorption regime. In the case of the IGM, even minute traces ($\sim$ few per cent) of intervening neutral hydrogen along the line-of-sight are capable of scattering Lyα photons away from the observer.

At $z \lesssim 3$, the IGM is on average very highly ionised except in dense self shielded clumps (e.g. Fan et al. 2006; McGreer et al. 2015; Planck Collaboration XIII 2016; Planck Collaboration XLVI 2016; Planck Collaboration XLVII 2016). Typically, at these redshifts the resonant scattering and absorption of Lyα photons occurs only due to diffuse amounts of neutral hydrogen which appear blueward of Lyα ($\lambda < \lambda_{\rm Ly\alpha}$) as a series of discrete narrow absorption features,

⋆ E-mail: bradley.greig@sns.it (BG),

© 0000 RAS

known as the Lyα forest (Rauch 1998). More problematic for measuring the intrinsic Lyα emission line are larger neutral column density absorbers such as Lyman-limit systems and damped Lyα absorbers (DLAs). In the case of DLAs, with column densities of $N_{\rm H\,I} > 10^{20}{\rm cm}^{-2}$ not only do these systems lead to completely saturated absorption, but they are also sufficiently dense to allow Lyα absorption within the Lorentzian wings. If these DLAs are located sufficiently close to the source QSO, absorption in the wings can significantly affect the intrinsic Lyα emission profile. Furthermore, associated strong absorbers intrinsic to the host QSO environment itself can additionally impact the Lyα line emission (Shen & Ménard 2012).

At $z > 6$, the IGM becomes increasingly neutral, and the Lyα forest gives way to completely dark (absorbed) patches (e.g. Barkana 2002; Gallerani et al. 2008; Mesinger 2010; McGreer et al. 2015). Once the IGM itself obtains a sufficiently large column density, it too begins to absorb Lyα in the Lorentzian wings of the scattering cross-section, referred to as IGM damping wing absorption (Miralda-Escudé 1998). In principle, through the detection of the IGM damping wing imprint from the Lyα emission spectra of high-z QSOs one is able to provide a direct measurement of the neutral fraction of the IGM during the reionisation epoch. Importantly, this approach requires an estimate of the intrinsic Lyα emission line in order to determine the contribution from the smooth component absorption of the IGM damping wing. Using this technique, limits on the IGM neutral fraction at $z > 6$ have been obtained (e.g. Mesinger & Haiman 2007; Bolton et al. 2011; Schroeder et al. 2013).

In the absence of a viable method to recover the intrinsic Lyα emission profile, one can attempt to use an emission template constructed from a QSO composite spectrum. Numerous composite spectra exist, constructed by averaging over a vastly different number of QSOs (e.g. Francis et al. 1991; Brotherton et al. 2001; Vanden Berk et al. 2001; Telfer et al. 2002; Shull et al. 2012; Stevans et al. 2014; Harris et al. 2016). However, by construction, these composite spectra only describe the average properties of QSOs, not the intrinsic variations of individual QSOs. Not accounting for these intrinsic properties can bias estimates of the IGM damping wing (Mortlock et al. 2011; Bosman & Becker 2015).

An alternative method is to reconstruct a template using a principle component analysis (PCA) (e.g. Boroson & Green 1992; Francis et al. 1992; Suzuki et al. 2005; Suzuki 2006; Lee & Spergel 2011; Pâris et al. 2011). An improvement over the use of a composite spectrum, this approach aims to use the minimal subset of eigenvectors to characterise the QSO emission profile. For example, Francis et al. (1992) find 3 eigenvectors are sufficient to describe 75 per cent of the observed profile variation, and 95 per cent if 10 eigenvectors are used. Beyond 10, the eigenvectors obtain little information and become rather noisy (Suzuki et al. 2005). However, these eigenvectors are extracted from a fit to the original QSO spectrum. Therefore, given that the Lyα line cannot be directly recovered at high redshifts, in order to obtain an estimate of the intrinsic Lyα profile some form of reconstruction/extrapolation of the relevant eigenvectors that describe Lyα would be required.

Conveniently, studies of these QSO composites and PCA approaches have revealed a wealth of information regarding correlations amongst various emission lines and other observable properties of the source QSO. For example, the first eigenvector of Boroson & Green (1992) showed that in the Hβ region a strong anti-correlation existed between [O III] and Fe II. Furthermore, Hewett & Wild (2010) performed an in-depth exploration to establish relationships between emission line profiles and the source systemic redshift to reduce the scatter and biases in redshift determination. More directly relevant for this work, both Shang et al. (2007) and Kramer & Haiman (2009) observe a strong correlation between the Lyα and C IV peak blue shifts.

Motivated by the existence of correlations amongst the various emission lines properties, and the lack of either a robust physical model of quasar emission regions (e.g. Baldwin et al. 1995) or a method to recover the intrinsic Lyα profile within high-z or heavily Lyα obscured QSOs, we propose a new method to reconstruct the intrinsic Lyα emission line profile for any high-z or Lyα obscured QSO. In this work, we develop our reconstruction method using QSOs selected from the Baryon Oscillation Spectroscopic Survey (BOSS; Dawson et al. 2013), a component of SDSS-III (Eisenstein et al. 2011). In summary, our Lyα reconstruction method is as follows:

- Perform a Monte Carlo Markov Chain (MCMC) fit to a subset of emission lines for each of our selected QSOs.
- Construct a covariance matrix which describes all correlations amongst the QSO emission lines.
- Assume a N-dimensional Gaussian likelihood function to describe the covariance matrix.
- MCMC fit a high-z or Lyα obscured QSO characterising all lines except Lyα.
- Use the recovered emission line information from the QSO to statistically characterise the intrinsic Lyα line.

The remainder of this paper is organised as follows. In Section 2, we discuss the observational data used within this work, and the selection criteria we apply to construct our sample of QSOs. In Section 3 we describe our MCMC fitting procedure, and how we model the QSO continuum, emission line features and other components to aid the estimation of the observed QSO flux. With all the data obtained from fitting our entire QSO sample, in Section 4 we construct our covariance matrix, and discuss the major correlations and recovered features. In Section 5 we outline our Lyα reconstruction method, and highlight the performance of this approach. Following this, in Section 6 we provide a discussion of the potential applications for both the MCMC fitting algorithm and the Lyα reconstruction pipeline. Finally, in Section 7 we finish with our closing remarks.

2 DATA

In this work, we select our QSOs from Data Release 12 (DR12) (Alam et al. 2015) of the large-scale SDSS-III observational programme BOSS (Dawson et al. 2013). The full details of the SDSS telescope are available in Gunn et al. (2006) and the details of the upgraded SDSS/BOSS spectrographs may be obtained in Smee et al. (2013). For reference, the wavelength coverage of the BOSS spectrograph is 3,600Å $< \lambda <$ 10,400Å in the observed frame, with a resolution of $R \sim 2000 - 2500$ corresponding to pixel resolution of $\sim 120 - 150$ km/s. The QSO target selection for BOSS

consisted of a variety of schemes, including colour and variability selection, counterparts to radio and X-ray sources, and previously known QSOs (Bovy et al. 2011; Kirkpatrick et al. 2011; Ross et al. 2012). The candidate QSO spectra were then visually inspected following the procedure outlined in Pâris et al. (2016), with updated lists of confirmed QSOs for DR12 available online[1]. Furthermore, we use the publicly available flux calibration model of Margala et al. (2016) to perform the spectrophotometric corrections of the DR12 QSO fluxes.

In total, DR12, contains 294,512 uniquely identified QSOs, 158,917 of which are observed within the redshift range $2.15 < z < 3.5$. The principal science goal of the BOSS observational programme was the detection of the baryon acoustic oscillation scale from the Lyα forest. To perform this, only relatively low to moderate S/N QSOs are required (e.g. White 2003; McDonald & Eisenstein 2007). However, within this work, we aim to statistically characterise the correlations between numerous QSO emission lines, which require much higher S/N.

To construct the QSO sample used within this work, we preferentially selected QSOs with a median S/N across all filters (*ugriz*) of S/N > 15 ('snMedian > 15'). This choice is arbitrary, and in principle we could go lower, however for decreasing S/N the weaker emission lines become more difficult to differentiate from the noise, reducing potential correlations. We selected QSOs containing broad-line emission using the 'BROADLINE' flag and removed all sources visually confirmed to contain broad absorption lines (BALs) by using the 'QSO' selection flag. Furthermore, only QSOs with 'ZWARNING' set to zero are retained, where the redshifts were recovered with high confidence from the BOSS pipeline. Finally, we restrict our QSO redshift range to $2.08 < z < 2.5$[2], selecting QSOs by their BOSS pipeline redshift[3] (e.g. Bolton et al. 2012; see Appendix A for a more in-depth discussion on our adopted choice of BOSS redshift). Following these selection cuts, we recover a total QSO sample of 3,862[4].

Though the total QSO number of 3,862 appears small, it should be more than sufficient to elucidate any statistically significant correlations amongst the emission line parameters[5]. Furthermore, to properly assess the MCMC fitting to be discussed in Section 3 a detailed visual inspection will be required, therefore it is preferential to restrict the sample size. In selecting our redshift range of $2.08 < z < 2.5$, we are assuming there is no strong variation in the emission line profiles of QSOs across the age of the Universe. That is, the covariance matrix recovered from this QSO sample will always be representative of the underlying QSO population. In Becker et al. (2013), these authors constructed 26 QSO composite spectra between rest-frame $1040 < \lambda < 1550$Å across the redshift range of $2 < z < 5$. They found for Lyα and the emission lines redward, no significant variation with redshift, lending confidence to our assumption.

Note, throughout this work, we do not deredden our QSOs to account for interstellar dust extinction within the Milky Way (e.g. Fitzpatrick 1999). The dust extinction curve varies strongly in the UV and ultimately will impact the Lyα line more than the C IV line. However, this will not greatly impact our results as we are attempting to reconstruct the Lyα emission line using a covariance matrix of correlations from other emission lines. Applying the extinction curve will only impact slightly on the peak of the emission lines, not the other characteristics of the line profile. When highlighting the performance of the reconstruction process in Section 5, we see that this slight overestimate should be well within the errors of the reconstructed profile.

3 MCMC QSO FITTING

We now introduce our MCMC fitting procedure. In Section 3.1, we outline and justify the selection of the emission lines considered. In Section 3.2 we then detail the construction of the QSO template. In order to improve the ability to recover the intrinsic emission line profile and QSO continuum we outline the treatment of 'absorption' features in Section 3.3. In Section 3.4 the iterative MCMC fitting procedure is discussed and an example is presented. Finally, in Section 3.5 we perform a visual quality assessment of our entire sample of fit QSOs to remove contaminants that could impact the observed emission line correlations.

3.1 Emission line selection

Owing to the limited wavelength coverage of BOSS, the selection of available emission lines to be used is limited. Due to the complexity in observing high-z ($z \gtrsim 6$) QSOs, where Lyα is redshifted into the near infrared (IR), the emission lines we use from the BOSS sample needs to be consistent with what is detectable with near-IR instruments such as Keck/MOSFIRE (McLean et al. 2010, 2012) and VLT/X-Shooter (Vernet et al. 2011).

The strongest emission lines, especially Lyα, C IV and Mg II, are known to contain a broad and narrow component (e.g. Wills et al. 1993; Baldwin et al. 1996; Vanden Berk et al. 2001; Shang et al. 2007; Kramer & Haiman 2009; Shen et al. 2011). These different components are thought to arise from moving clouds of gas above and below the accretion disk known as the broad and narrow line regions.

[1] http://www.sdss.org/dr12/algorithms/boss-dr12-quasar-catalog/

[2] This choice of redshift range was a trade off between the wavelength coverage of the spectrograph and requiring that we were sufficiently blueward of rest-frame Lyα to characterise the line profile (arbitrarily chosen to be 1180Å, corresponding to $\sim$ 3600Å at $z = 2.08$) while also including the Mg II emission line ($\lambda = 2798.75$Å, corresponding to $\lambda = 10,000$Å at $z = 2.5$).

[3] Note that the r.m.s of the quasar redshift error distribution for the BOSS pipeline redshift from DR9 was determined to be 550 km/s with a mean offset of the quasar redshift (bias) of ~ -150 km/s (Font-Ribera et al. 2013).

[4] We computed the absolute AB magnitude from the QSO continuum at 1450Å (M_{1450}) for this sample of QSOs and recovered a median of $M_{1450} = -26.1$ with an interquartile range of 0.7.

[5] The errors and biases when attempting to accurately measure the covariance matrix of a set of parameters ($N_{\rm par}$) from a total number of sources ($N_{\rm src}$) scales approximately as $N_{\rm par}/N_{\rm src}$ (e.g Dodelson & Schneider 2013; Taylor & Joachimi 2014; Petri et al. 2016). Within this work, we ultimately construct an 18 parameter covariance matrix from a subsample of 1673 QSOs, corresponding to an expected error of ~ 1 per cent.

Note that, for our purposes, the physical origins of these components is irrelevant; a double gaussian merely provides us a flexible basis set in which to characterise the line profile. Throughout this work, we will approximate all emission lines to have a Gaussian profile. While emission line profiles are known to be Lorentzian (e.g. Peebles 1993), the fitting of a Lorentzian profile is complicated by the requisite identification of the broader wings relative to the uncertainty in the QSO continuum[6]. Ultimately, this is a subtle difference as we are only interested in characterising the total line profile.

In order to decide which strong emission line profiles should be fit with a single or double component Gaussian[7], we perform a simple test analysing the Bayes information criteria (BIC; Schwarz 1978; Liddle 2004) of each of the individual line profiles in Appendix B. Below we summarise our findings:

- Lyα: The strongest observed UV emission line. In order to better characterise the fit to the line profile, a two component Gaussian is preferred. This is well known and consistent with other works.
- Si IV +O IV]: This line complex consists of both the Si IV (λ1396.76Å) and O IV] (λ1402.06Å) line centres. Unfortunately, these individual line profiles are intrinsically broad, preventing our ability to distinguish between them. Therefore these two lines appear as a single, blended line. In Appendix B we find a single Gaussian component is sufficient to characterise the line profile.
- C IV: This line profile is a doublet ($\lambda\lambda$1548.20, 1550.78Å). Again, both components being intrinsically broad prevents individual detection, therefore the line profile is observed as a single, strong line. For the C IV line, we find a two component Gaussian to be preferable to characterise the line.
- C III]: For the C III] line (λ1908.73Å) we find a single component Gaussian to be sufficient to characterise the line. For almost all of the BOSS QSOs in our sample, the nearby, weak Si III] line (λ1892.03Å), which would appear on top of the much broader C III] line is not resolvable[8]. Furthermore, the C III] line may also be contaminated by a continuum of low ionisation Fe lines, however, unlike Mg II (see next) the relative impact on the C III] line is minor.
- Mg II: While the Mg II line (λ2798.75Å) is observable within the wavelength range of our BOSS QSO sample, we choose not to fit this emission profile. This is because of the contamination from the Fe pseudo-continuum. Though Fe emission templates exist (e.g. Vestergaard & Wilkes 2001), the fitting procedure is complicated. One must simultaneously fit several spectral regions in order to calibrate the true flux level of the Fe pseudo-continuum in order to remove it and obtain an estimate of the Mg II emission line. Since we aim to simultaneously fit the full QSO spectrum, this approach would add degeneracies to the model (e.g. between the true continuum, and the Fe pseudo-continuum). While this can be overcome, it requires first fitting the QSO continuum blueward of Mg II, before fitting the Fe pseudo-continuum, which reduces the flexibility of our model.

In addition to these lines above, we also consider several other lines all modelled by a single Gaussian. These include: (i) N V ($\lambda\lambda$1238.8, 1242.8Å), which is an important line for characterising the Lyα line profile as it can be degenerate with the broad Lyα line component (ii) Si II (λ1262.59Å), (iii) the O I/Si II blended complex (λ1304.35, λ1306.82Å), (iv) C II (λ1335.30Å), (v) He II (λ1640.42Å) (vi) O III (λ1663.48Å) and (vii) Al III (λ1857.40Å).

3.2 Continuum and emission line template

In this work, we fit the rest-frame wavelength range (1180Å $< \lambda <$ 2300Å) with a single power-law for the QSO continuum,

$$f_\lambda = f_{1450} \left(\frac{\lambda}{1450\text{Å}} \right)^{\alpha_\lambda} \text{erg cm}^{-2}\,\text{s}^{-1}\,\text{Å}^{-1}, \quad (1)$$

where α_λ describes the spectral slope of the continuum and f_{1450} is the normalisation of the QSO flux which we choose to measure at 1450Å. While we normalise the QSO flux at 1450Å, we allow this quantity to vary within our MCMC fitting algorithm (by adding a small, variable perturbation). While this quantity does not depart greatly from the original normalised value, it allows us to compensate for situations where the nearby region around 1450Å might be impacted by line absorption or a noise feature from the spectrograph.

Secondary (broken) power-law continua have been fit to QSO spectra at $\lambda >$ 4500Å for an independent red continuum slope (e.g. Vanden Berk et al. 2001; Shang et al. 2007), however, this is beyond our QSO sample wavelength coverage. Typically, a broken power-law continuum is also adopted blueward of Lyα, visible in low-z HST spectra (e.g Telfer et al. 2002; Shull et al. 2012). In this work, we do not consider a different slope near Lyα as we only fit down

[6] In practice, thermal (or other) broadening will make the Lorentzian wings irrelevant in the case of AGN emission lines. For even modest thermal broadening of $\sim$ 10 km/s only the Gaussian core can be detected (see e.g. figure 4 of Dijkstra 2014).

[7] In principle, we could allow for $N > 2$ components to our line fitting procedure. For example, several authors have fitted three Gaussian components to the C IV emission line profile (e.g. Dietrich et al. 2002; Shen et al. 2008, 2011). However, in this work, we attempt to avoid adding too much complexity to the fitting method.

[8] In principle, by ignoring the Si III] feature, the centroid of the measured C III] line can be shifted bluer than its true value. Quantitatively, we determine the extent of this blueshift by fitting all QSOs within our sample that have a clear, discernible Si III] feature. We find that by excluding Si III] the average recovered C III] emission line centre is blue shifted by $\sim$ 310 km/s (with the associated line width broadened by $\sim$ 315 km/s). This sample, however, only contains 57 (out of 1673) QSOs, corresponding to only $\sim$ 3.4 per cent of the full 'good' sample affected by this blueshift. For the vast majority of QSOs in our sample the S/N around C III] is insufficient to tease out the individual emission component of the Si III] feature. Therefore, owing to the small resultant blueshift in the line centre estimate (compared to the average dispersion of $\sim$ 550 km/s in the redshift estimate of the chosen BOSS pipeline redshift) and the corresponding low number of sources potentially affected by this systematic shift, we deem the exclusion of Si III] a valid simplification. As a final point, owing to the weak total flux expected in the Si III] emission line, including this in the MCMC fitting for the full QSO sample (where Si III] is not clearly present) would cause a strong degeneracy with the C III] line unless extremely strong priors on the line shape are imposed.

to 1180Å, therefore there is insufficient information to include a secondary component. By only considering a single power-law through the Lyα line profile, we may bias our results slightly on fitting the Lyα broad component. However, provided we are consistent in our usage of this QSO continuum between this fitting approach and the reconstruction method, this should not impact our results.

As mentioned in the previous section, we model each emission line component with a Gaussian profile. Following Kramer & Haiman (2009) the total flux for each component can be defined as,

$$F_i = a_i \exp\left[-\frac{(\lambda-\mu_i)^2}{2\sigma_i^2}\right], \quad (2)$$

where a_i describes the amplitude of the line peak, μ_i is the location of the line centre in Å, σ_i is the width of the line in Å and the subscript 'i' denotes the specific line species (e.g. Lyα). Note that within this work, the peak amplitude is always normalised by the continuum flux at 1450Å, f_{1450}, therefore it is always a dimensionless quantity. More intuitively, the line centre location, μ_i, can be written in terms of a velocity offset relative to the systemic line centre,

$$v_{\mathrm{shift},i} = c\,\frac{\mu_i-\lambda_i}{\lambda_i}\ \mathrm{km/s}, \quad (3)$$

and the line width can be expressed as

$$\sigma_i = \lambda_i\left(\frac{\mathrm{LW}}{c}\right)\ \text{Å}, \quad (4)$$

where LW is the line width measured in km/s. Throughout Equations 2-4, both λ and λ_i are measured in the rest frame.

Each Gaussian line component can therefore be fully described by its three component parameters, the line width, peak amplitude and velocity offset. In total, we fit each QSO with two continuum parameters, two double component Gaussians (Lyα and C IV), and 9 single component Gaussian profiles, resulting in a total of 41 continuum and emission line parameters.

3.3 Identifying absorption features

Intervening diffuse neutral hydrogen or larger column density absorbers (smaller than DLAs) along the line of sight can produce narrow absorption features. Furthermore, metal pollution in stronger absorption systems can additionally result in narrow absorption features appearing in the observed profiles of the emission lines.

These narrow features, if not measured and accounted for, can artificially bias the shape and peak amplitude of the emission line components. Therefore, in this section we outline our approach for identifying these features in a clean and automated manner:

- Identify all flux pixels that are a local minima within a 2Å region surrounding the central pixel in question. This choice of 2Å is arbitrary, but is selected to be sufficiently broad to ignore features that might arise from noise fluctuations.
- Construct a horizontal line of constant flux that begins from the global flux minimum.
- Incrementally increase this line of constant flux, recording the depth and width of each absorption feature enclosed by the line of constant flux.
- If the depth becomes larger than 3σ (5σ in the vicinity of Lyα) of the observed error in the flux at that pixel and the candidate absorption line has remained isolated (i.e. not overlapped with a nearby feature) it is then classified as an absorption line.

Our adopted choice of 3σ (5σ near Lyα) arises after rigorously testing this pipeline against the absorption features that were visually identified, until the vast majority of the absorption lines could be robustly located by this procedure. Across the full QSO sample, we found a broad range in the number of features identified within each individual QSO, from only a few up to ~ 40. For each identified absorption feature, a single Gaussian (described by its own three parameters) is assigned and is simultaneously fit with the continuum and emission line profiles outlined previously. Therefore, the total number of parameters to be fit per QSO varies.

3.4 MCMC sampling the QSO template

We fit each QSO within a Bayesian MCMC framework, using the χ^2 likelihood function to determine the maximum likelihood fit to the QSO spectrum. This choice enables us to fully characterise any potential model degeneracies between our model parameters, while also providing the individual probability distribution functions (PDFs) for each model parameter. In this work, we utilise the publicly available MCMC python code CosmoHammer (Akeret et al. 2013) built upon EMCEE (Foreman-Mackey et al. 2013) which is based on the affine invariant MCMC sampler (Goodman & Weare 2010).

Assuming flat priors across our > 50 model parameters and fitting the full high-resolution BOSS spectrum within the wavelength range (1180Å $< \lambda <$ 2300Å) simultaneously is computationally inefficient. Instead, we perform an iterative procedure to boost the computational efficiency which we outline below:

- Normalise the full QSO spectrum at 1450Å and then fit the two continuum parameters, f_{1450} and α_λ within a set of selected wavelength ranges of the full spectrum which are minimally contaminated by emission lines. We choose to fit the QSO continuum within the regions [1275, 1295], [1315, 1330], [1351, 1362], [1452, 1520], [1680, 1735], [1786, 1834], [1970, 2040] and [2148, 2243] Å.
- Break the full QSO spectrum up into wavelength regions centred around the emission lines. We choose [1180, 1350] Å centred on Lyα, [1350, 1450] Å centred on Si IV +O IV], [1450, 1700] Å centred on C IV and [1700, 1960] Å centred on C III].
- Within each of these four regions, we take the continuum estimated from above, and then fit all emission lines and any absorption features which fall within the respective wavelength ranges. For each of these regions we then perform an MCMC fit.
- From the individual PDFs for each parameter we construct a flat prior across a much narrower allowed range, driven by the width of the individual distributions from the fitting above.
- Finally, we fit the entire QSO spectrum using the entire model parameter set and recover a maximum likelihood model which describes the full spectrum.

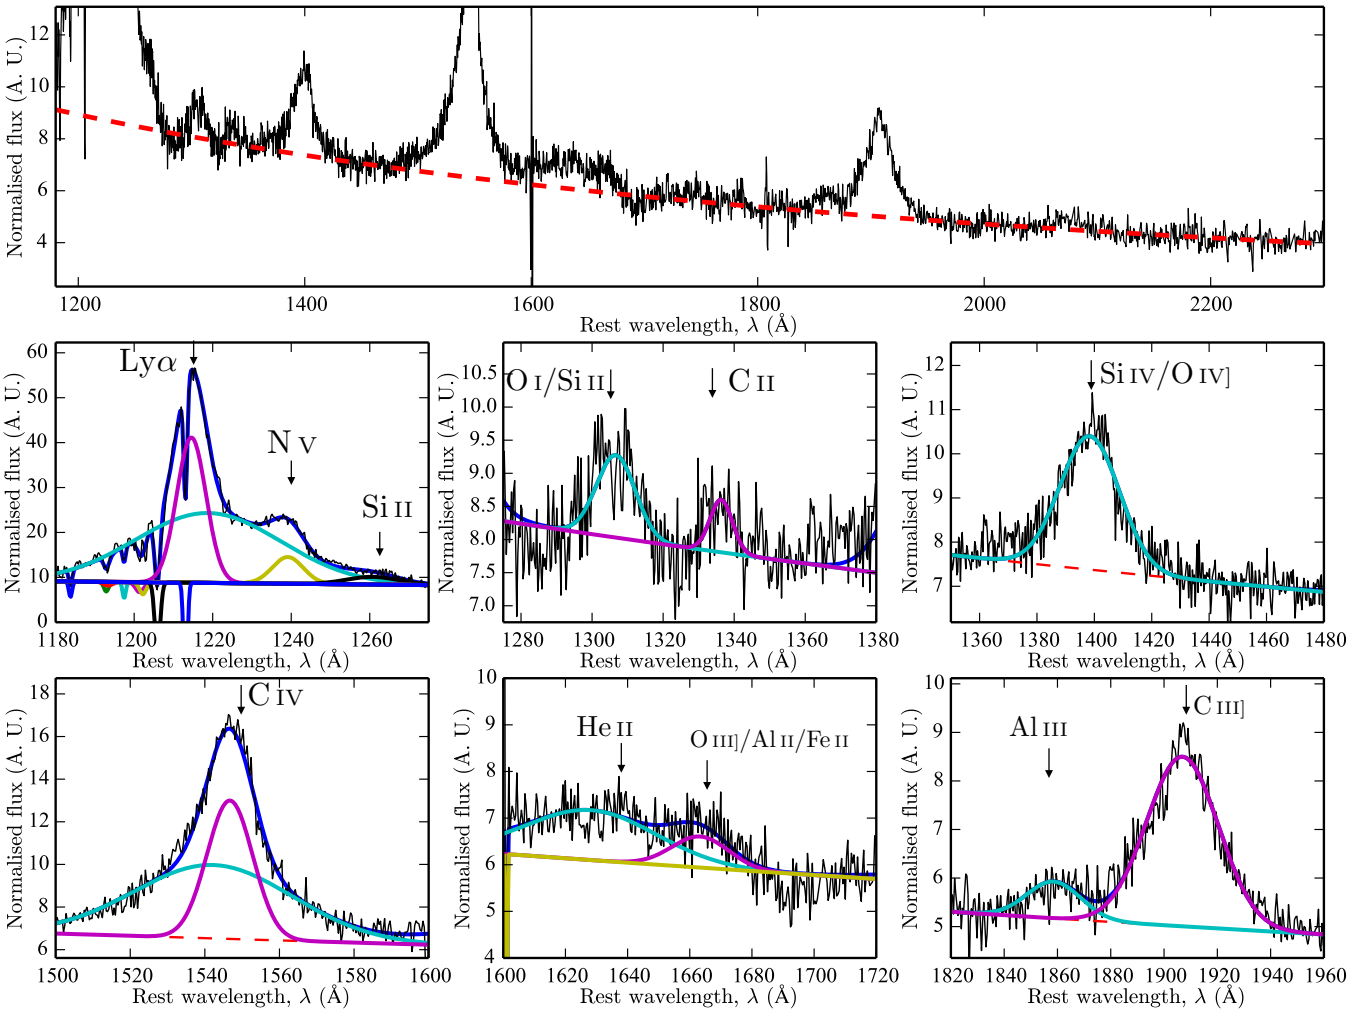

Figure 1. An example, zoomed in figure of the MCMC QSO template fitting of a BOSS (SDSS-III) spectrum ('spec-4185-55469-0076' a $z = 2.478$ QSO). For all spectra we translate the flux into arbitrary units (1 A. U. $= 10^{-17}\,\mathrm{erg\,cm^{-2}\,s^{-1}\,\AA^{-1}}$). *Top:* The full QSO spectrum used within our fitting procedure as a function of the rest frame wavelength, λ. The red dashed curve corresponds to the two parameter continuum (see Equation 1). *Middle left:* The two component fit to the Lyα emission profile highlighted by the broad (cyan) and narrow (magenta) components and a single Gaussian profiles for N V (yellow) and Si II (black). A series of 'absorption' features were identified and fit to the spectrum (shown below the continuum level). *Middle centre*: Two low ionisation lines, O I (cyan) and C II (magenta). *Middle right:* Single component fit to the Si IV +O IV] doublet. *Bottom left:* Double component fit to C IV highlighted by the broad (cyan) and narrow (magenta) components. *Bottom centre:* Low ionisation lines, He II (cyan) and O III (magenta). *Bottom right:* Single component fit to C III] (magenta) and single Gaussian Al III component (cyan). The maximum likelihood fit to this spectrum was $\chi^2 = 3110$, for which we recovered a $\chi^2_{\rm red} = 1.1$ (with 2828 degrees of freedom, 2899 bins and 71 free parameters).

In Figure 1, we provide an example of one of the BOSS QSOs from our full sample. In this figure we provide zoomed in panels of both the QSO continuum (red dashed curve, top panel) and the various emission lines which are simultaneously fit within our MCMC approach across all other panels. This figure highlights that any notable absorption features have been identified by our pipeline and accurately characterised (e.g. the Lyα line profile in the left panel of the middle row). From Figure 1 we note that this approach is able to fit the full spectrum. The maximum likelihood we obtained was $\chi^2 = 3110$, for which we had 2899 bins (from the raw spectrum) and 71 free parameters, corresponding to a reduced χ^2 of 1.1. For reference, this took $\sim$1 hour on a single processing core, which can be rapidly improved if a binned spectrum is used.

3.5 QSO fitting quality assessment

After fitting the entire sample with our full MCMC QSO fitting pipeline, we are in a position to construct a higher fidelity sample of QSOs for constructing our covariance matrix and investigating the correlations amongst the emission line parameters. This is performed by visually inspecting our QSO sample, applying a simple selection criteria. Following the completion of this selection process, we produce two separate QSO samples, one classified as 'good' and another classified as 'conservative', the details of which we discuss below (in Appendix C we provide a few select examples to visually highlight this subjective process). The criteria are outlined as follows:

- We remove all QSOs with a poor characterisation of the continuum. These include QSOs with a positive spectral

index, or a clear departure from a single power-law continuum. Only a handful of QSOs exhibit this behaviour. See Figure C1 for some examples.

- We further remove any QSOs which have either (i) missing sections of flux that overlap with any of the emission lines (ii) a sufficient number of absorption features which cause a loss of confidence in the fitting of the Lyα profile (iii) either an intervening dense neutral absorber or sufficiently strong/broad absorption blueward of line centre which can impact the broad component of either the Lyα or the C IV lines (iv) a Lyα line region that is not well fit or characterised by our double component Gaussian, which might arise from numerous absorption profiles or the lack of a prominent Lyα peak.
- Following the removal of these contaminants from our sample, we are left with 2,653 QSOs, which we call our 'conservative' sample[9]. These can still include QSOs which might have absorption features centred on the Lyα line centre, or absorption line complexes which might contaminate large sections of the Lyα line profile (see Figure C3 for examples).
- We then apply a secondary criteria, which preferentially selects QSOs which contain (i) fewer absorption features and (ii) no absorption features on the line centre. This results in a final sample of 1673 QSOs, which we refer to as our 'good' sample (see Figure C4 for examples).

The major difference between the 'good' and 'conservative' sample is that the 'conservative' sample will contain QSO spectra for which we have less confidence in either the identification and fitting of one, or several of the emission lines due to the presence of absorption features. Ultimately, if our claim that any correlations amongst the emission lines are a universal property of the QSOs then the covariance matrices of the two should be almost identical, with the 'conservative' sample containing additional scatter (slightly weaker correlations). In the next section, we construct our covariance matrices and investigate this further.

4 DATA SAMPLE COVARIANCE

With the refined, quality assessed QSO spectra from the previous section, we now construct our covariance matrix to characterise all correlations amongst the various emission lines.

4.1 The covariance matrix

In performing the quality assessment of our QSO spectra, we note that a number of the weaker emission lines are not always well characterised or resolved. This becomes more prevalent for the QSOs nearer to our sample limit of S/N = 15, which correspond to the highest density of QSOs in our sample. It would be interesting to investigate correlations between the strong high ionisation lines and the weaker low ionisation lines, as well as correlations amongst these two classifications. However, since these lines are not always readily available in our QSOs we refrain from doing so. Regardless, by still attempting to fit these weaker lines we retain the flexibility of the MCMC approach, and more importantly this enables the QSO continuum to be estimated to a higher accuracy.

In light of this, we construct our covariance matrix from a subset of these emission lines: the most prominent emission lines that should always be resolvable in a lower S/N or lower resolution QSO spectrum (as in the near-IR for $z > 6$ QSOs). The lines we identify are Lyα, Si IV +O IV], C IV and C III]. Since for Lyα and C IV we allow a broad and narrow component, we recover an 18×18 covariance matrix. We choose to exclude the N V line from our covariance matrix for two reasons: (i) a clear, identifiable N V line is not always present in our QSO spectra, therefore including it would artificially reduce any visible correlation and (ii) the ultimate goal of this work is the reconstruction of the Lyα line whereby including the N V line to the reconstruction process would only increases the complexity (the N V line should in principle be recoverable from the high-z or Lyα obscured QSO).

In Figure 2, we present the correlation coefficient matrix, which is obtained from the full covariance matrix from our 'good' sample of 1673 QSOs. In constructing this, we assume the standard format for the covariance matrix, $\mathbf{\Sigma}_{ij}$, given by,

$$\mathbf{\Sigma}_{ij} = \frac{1}{N-1}\sum_{i}^{N}(\mathbf{X}_i - \boldsymbol{\mu}_i)(\mathbf{X}_j - \boldsymbol{\mu}_j), \tag{5}$$

where $\mathbf{X}_i$ is the data vector for the full QSO sample and $\boldsymbol{\mu}$ is the mean data vector for the ith parameter. For this data vector we use the values from the parameter set which provide the maximum likelihood fit to each QSO. The correlation coefficient matrix, $\mathbf{R}_{ij}$, is then defined in the standard way,

$$\mathbf{R}_{ij} = \frac{\mathbf{C}_{ij}}{\sqrt{\mathbf{C}_{ii}\mathbf{C}_{jj}}}, \tag{6}$$

where each diagonal entry, $\mathbf{R}_{ij}$, corresponds to the correlation coefficient between the ith and jth model parameters.

Note, for the covariance matrix, we do not include the two continuum parameters. Firstly, for the parameters defining the emission line peak we define this parameter as the normalised peak value, where it is normalised by the continuum flux at f_{1450}. Therefore, if included, this would be completely degenerate with the peak amplitudes. Some correlation with continuum spectral index is expected, given that external (to the broad-line region) reddening will simultaneously weaken the bluer UV lines and redden the continuum spectral index. However, mild reddening is seen in only $\sim 10-20$ per cent of SDSS quasars (e.g. Richards et al. 2003; Hopkins et al. 2004) so this has little effect on the correlations (we find very weak correlations on the order of 10 per cent). Therefore we do not report them, as these provide no additional information with respect to the individual line correlations.

[9] Our naming convention of 'good' and 'conservative' refers to the tightness of constraints recovered from the reconstruction procedure. The 'conservative' sample, with less assumptions on data quality, produces slightly broader errors relative to the 'good' sample.

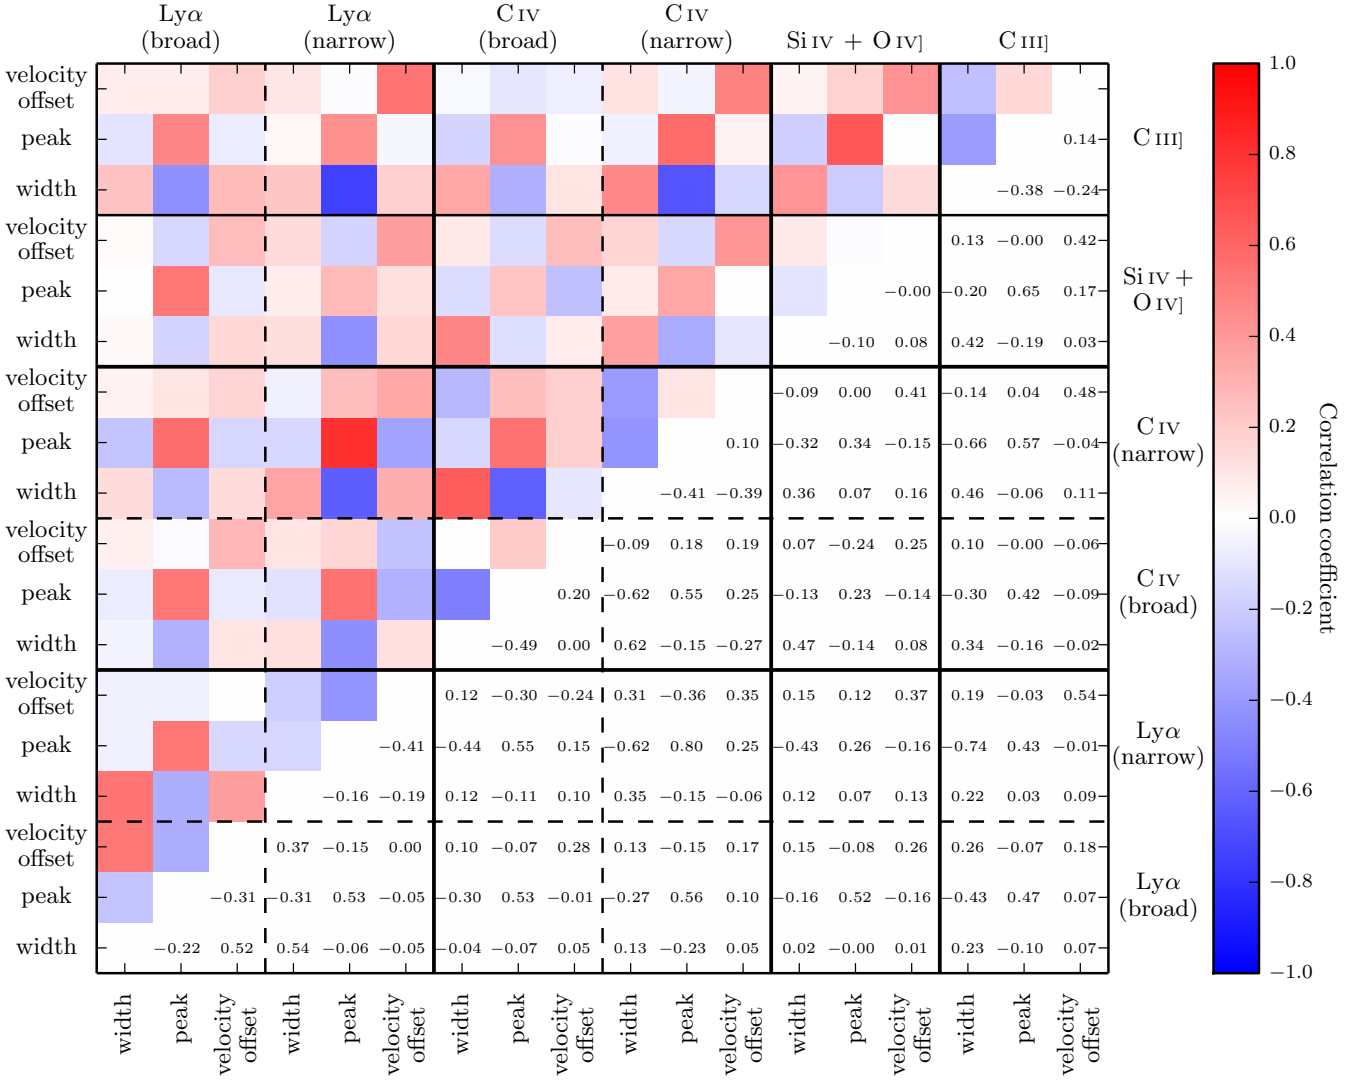

Figure 2. The correlation coefficient matrix (correlation coefficients listed in the lower half) constructed from the 'good' sample of 1673 QSOs. This 18×18 matrix contains the double component Gaussians of Lyα and C IV and the single component Gaussians for Si IV +O IV] and C III]. Each Gaussian component contains three parameters, its peak width, peak height and velocity offset from systemic (Lyα + C IV + Si IV +O IV] + C III] $= (2 \times 3) + (2 \times 3) + 3 + 3 = 18$ parameters).

4.2 Interpreting the covariance matrix

In order to aid the interpretation of the correlation matrix, we divide Figure 2 into the four emission line species (denoted by black curves), while the narrow and broad line species are further separated by black dashed lines. Positive correlations are represented by a decreasing (weakening) shading of red, with white representing no correlation and an increasing blue scaling denotes a strengthening of the anti-correlation. Using the upper half of the correlation matrix enables a faster analysis of the correlation patterns amongst the emission line parameters, and the lower half reports the numerical value of the correlation (or anti correlation). For the most part, each 3×3 sub-matrix returns the same correlations and anti-correlations amongst the peaks, widths and velocity offsets, but with varying degrees of strength.

Firstly, we note that with this correlation matrix it is straightforward to notice that all pairs of peak height parameters are positively correlated. This is to be expected as the peak amplitude is positively correlated with the QSO luminosity. This correlation between the peak heights is the strongest of the trends. Most importantly for this work, the correlations amongst the Lyα peak height parameters are the strongest. For example, we find the strongest correlation ($\rho = 0.8$) between the peak height of the Lyα narrow component and the associated peak height of the C IV narrow component. In Figure 3, we provide the 2D scatter plot for this strong correlation (central panel) and the 1D marginalised PDFs for the Lyα (top) and C IV (right). The green solid and dashed contours in the central panel denote the 68 and 95 per cent 2D marginalised joint likelihood contours, which describe the relative scatter amongst our sample of QSOs. In the 1D marginalised PDFs, the solid black curves are the histograms of the sample of QSOs, while the blue solid curves are an approximated 1D Gaussian for that associated parameter. Note in this figure, and in subsequent figures, these blue curves are not a fit directly to the raw data, rather instead they are approximations of a Gaussian PDF with 1σ scatter equivalent to the raw data.

Returning to the correlation matrix, we additionally re-

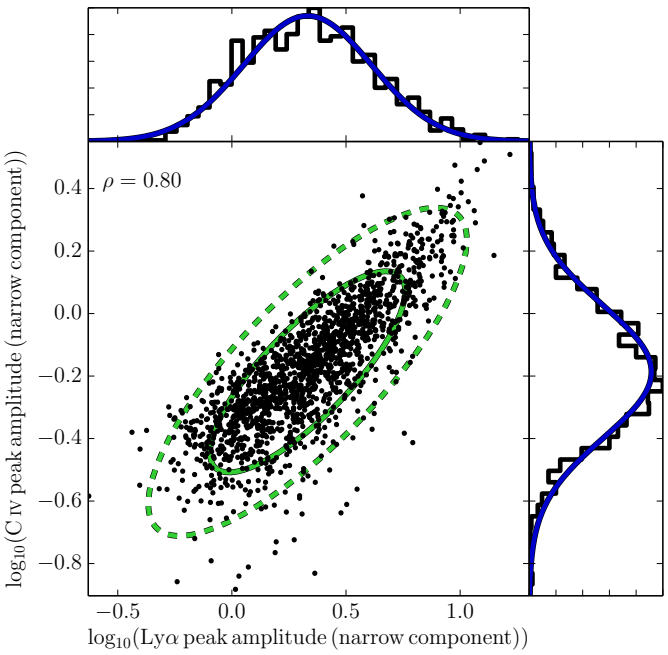

Figure 3. A 2D scatter plot of the correlation between the peak amplitudes of the narrow components of the Lyα and C IV emission lines for our 'good' sample of 1673 QSOs. We recover a correlation coefficient of $\rho = 0.80$, indicative of a strong, positive correlation. The green solid and dashed curves correspond to the 68 and 95 per cent 2D marginalised joint likelihood contours, respectively. The histograms (black curves) correspond to the recovered PDFs for each of the two parameters, while we also provide Gaussian curves (blue) representative of the scatter in the fit parameters (not a direct fit). Note that for the peak amplitude, we normalise by f_{1450}, therefore these are dimensionless.

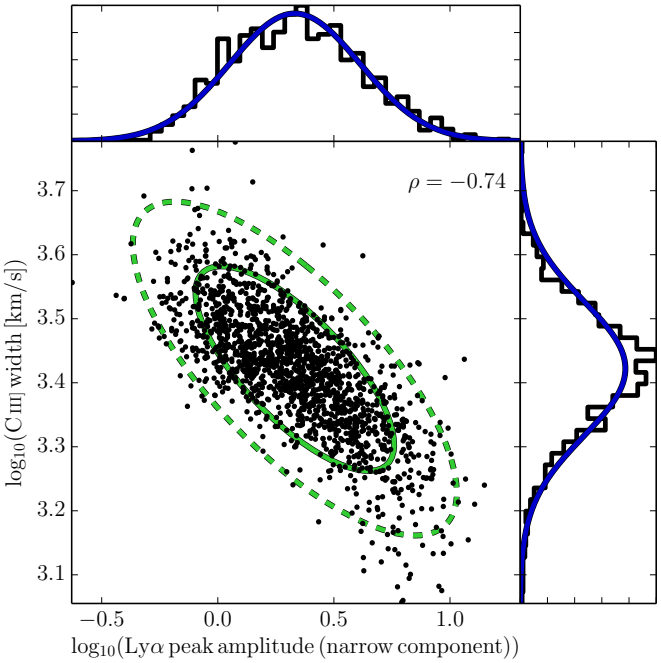

Figure 4. A 2D scatter plot of the correlation between the narrow component of the Lyα peak amplitude and line width of the C III] emission line for our 'good' sample of 1673 QSOs. We recover a correlation coefficient of $\rho = -0.74$, indicative of a strong anti-correlation. Histograms, blue curves and the green contours are as described in Figure 3. As in Figure 3, the peak amplitude is normalised by f_{1450} and is therefore dimensionless.

cover a relatively strong trend (anti-correlation) between the peak height and the line width. In Figure 4, we provide the strongest of these anti-correlations ($\rho = -0.74$) which is between the Lyα peak amplitude of the narrow line component and the width of the single component C III] line. We observe, that for an increasing peak amplitude for the Lyα narrow component, a decreasing of the C III] line width. This behaviour could be inferred as a result of the Baldwin effect (Baldwin 1977). For an increasing peak amplitude (i.e. QSO luminosity), we expect a weaker broad-line emission. Physically, a plausible scenario to describe this could be the over-ionisation of the inner broad-line region by the continuum in high luminosity QSOs, resulting in carbon being ionised into higher order species (i.e. no C IV/C III]). For these high luminosity QSOs, the high ionisation lines may then predominantly arise at larger radii where the velocity dispersion is lower, producing smaller line widths (e.g. Richards et al. 2011). Alternately, it could arise from systematics from our line fitting. As a result of the decreasing equivalent widths with increasing QSO luminosity, it could be that the single component Gaussian for the C III] does not characterise the line profile as accurately. With a less prominent broad-line component, the emission from the wings could be underestimated, producing a narrower C III] line.

We additionally recover a trend for a positive correlation between the widths of the various line species, though this is a relatively weaker trend. However, for the broad and narrow components for the same line species, we find more moderate correlations ($\rho = 0.54$ for Lyα and $\rho = 0.62$ for C IV). It is difficult to interpret these correlations though, as these broad and narrow lines are simultaneously fit, and therefore in principle could be degenerate.

Shang et al. (2007) investigated correlations amongst the Lyα, C IV and C III] lines for a significantly smaller sample of 22 QSOs. These authors recover moderate to strong correlations between a few of their emission line width full width half maxima (FWHMs). They find Lyα–C IV ($\rho = 0.81$) and C IV–C III] ($\rho = 0.53$). If we equate this FWHM as the width of the narrow component of our emission line profiles we find correlations between the same species of Lyα–C IV ($\rho = 0.35$) and C IV–C III] ($\rho = 0.46$)[10]. While we recover an equivalent correlation for the C IV–C III] lines, we find a significant discrepancy for the Lyα–C IV lines. This potentially can be explained as an artificially large correlation owing to their small number of objects (22 compared to our 1683 QSOs)[11]. Supporting this hypothesis, in Corbin & Boroson (1996) for a larger sample of 44 QSOs, this is reduced to $\rho = 0.68$ for Lyα–C IV.

Additionally, Kramer & Haiman (2009) recovered a moderate positive correlation ($\rho = 0.68$) between the veloc-

[10] By converting our fitted emission line parameters (both broad and narrow) into a line profile FWHM width, we find correlations similar in strength to those recovered purely from the narrow line component.

[11] Note that, if we adopt a systemic redshift measured purely from a single ionisation line (such as the Mg II redshift) rather than the pipeline redshift used throughout this work, we can recover a significantly stronger correlation of Lyα–C IV ($\rho = 0.7$).

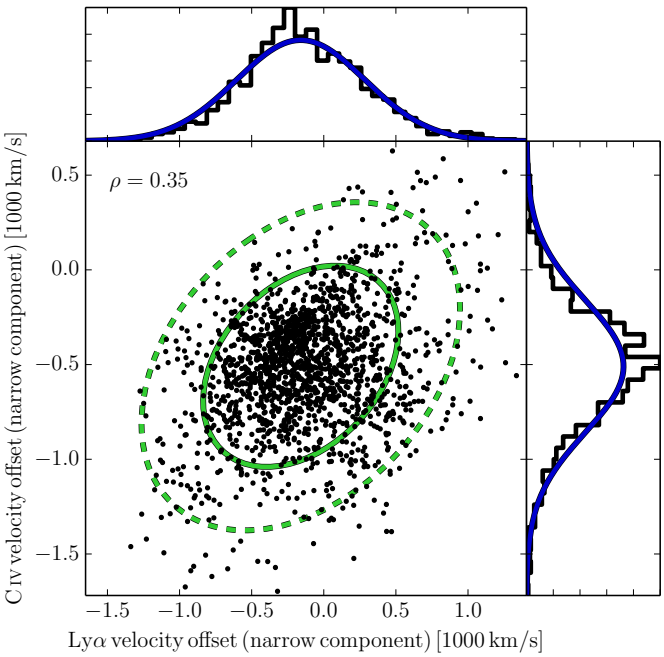

Figure 5. A 2D scatter plot of the correlation between the velocity offsets of the Lyα and C IV emission lines for our 'good' sample of 1673 QSOs. We recover a correlation coefficient of $\rho = 0.35$, indicative of a moderate correlation, which is counter to the much stronger correlation hinted at from the smaller QSO samples of Shang et al. (2007) and Kramer & Haiman (2009). Histograms, blue curves and the green contours are as described in Figure 3.

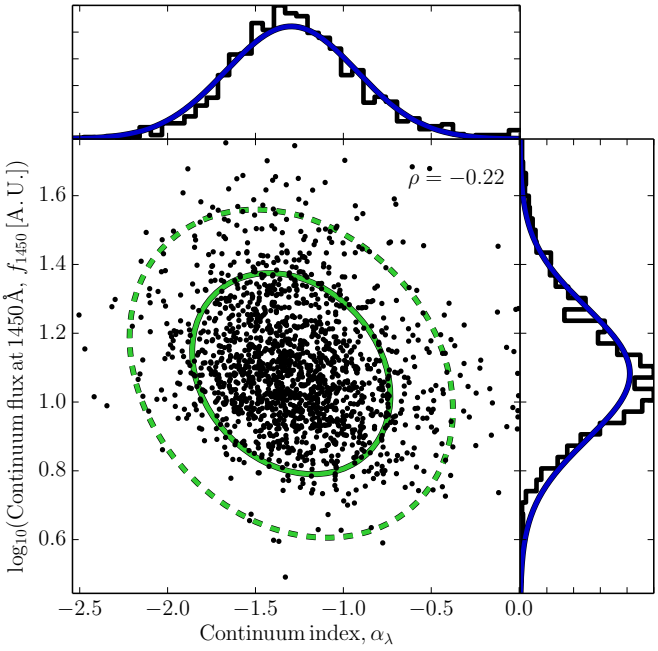

Figure 6. A 2D scatter plot highlighting the correlation between the two QSO continuum parameters: the continuum spectral index, α_λ and the continuum flux at 1450Å, f_{1450} (for which we define 1 A. U. $= 10^{-17}\,\mathrm{erg\,cm^{-2}\,s^{-1}\,\AA^{-1}}$), for our sample of 1673 QSOs (the 'good' sample). We recover a correlation coefficient of $\rho = -0.22$, indicative of a weak anti-correlation. Histograms, blue curves and the green contours are as described in Figure 3.

ity offsets for the Lyα–C IV components. For the same combination, Shang et al. (2007) found a stronger correlation of $\rho = 0.81$. In Figure 5, we provide our weaker results between the narrow components of the Lyα–C IV lines for which we recover a moderate correlation of $\rho = 0.35$. Once again, this lack of an equivalently strong correlation in our sample could arise purely from the fact that our sample contains two orders of magnitude more QSOs. Shang et al. (2007) equivalently quote correlations of $\rho = 0.45$ (Lyα–C III]) and $\rho = 0.39$ (C IV–C III]). Using the narrow lines for Lyα and C IV, we find similarly moderate, but slightly stronger correlations of $\rho = 0.54$ (Lyα–C III]) and $\rho = 0.48$ (C IV–C III]).

In the construction of the covariance matrix, and associated correlation matrix we have only used the maximum likelihood values, and ignored the relative errors for each parameter from the marginalised 1D PDFs. Our reasoning for this is that the amplitude of the error on the individual parameters is small with respect to the scatter that arises for each parameter across the full QSO sample. In principle however, we could perform an MCMC sampling of the full covariance matrix, allowing the means and all correlation coefficients to be free parameters. This would allow a more accurate characterisation of both the individual errors for each parameter and the total scatter across the full QSO sample, potentially tightening the correlations amongst the emission line parameters. However, this would require fitting $\frac{1}{2}N(N+1)$ parameters encompassing both the means of the individual parameters and all correlation coefficients (in our case, 171 free parameters).

4.3 Correlation of the QSO continuum parameters

Though the continuum parameters are not included in the covariance matrix in Figure 2, in Figure 6 we provide the 2D scatter between these two parameters. In our 'good' sample of QSOs, we recover only a relatively weak anti-correlation ($\rho = -0.22$) between the QSO continuum spectral index, α_λ and the normalisation at 1450 Å (f_{1450}). This weak anti-correlation likely arises due to dust reddening, which causes a drop in the normalisation, f_{1450} and shallower spectral slopes. While both of these QSO parameters are well characterised by a Gaussian, the respective scatter in each parameter is considerable.

For our sample of 'good' ('conservative') QSOs, we recover a median QSO spectral index of $\alpha_\lambda = -1.30 \pm 0.37$ (-1.28 ± 0.38). In contrast, Harris et al. (2016) construct a QSO composite spectrum with a wavelength coverage of (800Å $< \lambda <$ 3300Å) from the same BOSS DR12 sample. Using $\sim 100,000$ QSOs, these authors find a median spectral index for their QSO sample of $\alpha_\lambda = -1.46$, consistent with our results within the large scatter. However, these authors only fit the QSO continuum between 1440-1480Å and 2160-2230Å. Not fitting to the same spectral region likely results in a different recovered spectral slope. Furthermore, our estimate of $\alpha_\lambda = -1.30$ ($\alpha_\nu = -0.70$) is also consistent with the lower redshift samples of Scott et al. (2004) ($\alpha_\nu = -0.56^{+0.38}_{-0.28}$), Shull et al. (2012) ($\alpha_\nu = -0.68 \pm 0.14$) and Stevans et al. (2014) ($\alpha_\nu = -0.83 \pm 0.09$). Note, within all these works, considerable scatter in the spectral index is also prevalent.

4.4 Potential sample bias

In Sections 4.1 and 4.2, we presented the correlation matrix and discussions on the relative trends between the emission lines and their relative strengths. However, these results were drawn from our refined, quality assessed 'good' QSO sample. In order to guard against a potential bias which may have arisen following our specific selection process, we additionally construct the correlation matrix for our 'conservative' QSO sample. This 'conservative' sample contains $\sim$ 1000 additional QSOs which are defined to be less robust than required for our 'good' sample. Therefore, this 'conservative' sample should contain more scatter amongst the recovered parameters, which would notably degrade the strength of the correlations relative to the 'good' sample if we have artificially biased our results.

In Figure 7 we provide the matrix of the relative difference in the correlation coefficients between the 'good' and 'conservative' QSO samples. Here, we show the amplitude of the change in correlation coefficient between the two QSO samples. A positive (red) difference is indicative of the 'good' sample having a stronger correlation (either positive or anti-correlation), while a negative (blue) difference is indicative of the 'conservative' sample having a stronger correlation. Note, squares marked with a dot-dashed cross indicate a change between a positive and anti-correlation, which arises when the correlations are close to zero in either sample.

For the most part, the relative change in the correlation is minor, of the order of $|\Delta\rho| < 0.04$. More importantly, for any of the strong correlations and notable trends we discussed in the previous section, we observe no sizeable differences, with the 'good' sample providing on average slightly stronger correlations (as indicated by the prevalence of red squares). For example, the Lyα–C IV peak heights of the narrow component from our 'good' sample was $\rho = 0.80$, whereas for the 'conservative sample it is $\rho = 0.79$. Therefore, in the absence of any drastic differences in the strong and notable trends discussed previously between the two samples, it is clear that we have not biased our QSO covariance matrix in the construction of the 'good' sample. Given the relative amplitude of these differences between the two QSO samples is small, and tend to be slightly weaker for the 'conservative' sample, the reconstructed profile recovered from the covariance matrix should recover effectively the same best-fit Lyα profile, with slightly broader errors owing to the increased scatter (reduced correlations).

Equally, to be prudent, we considered several other ways to divide our QSO sample. First, we performed the same analysis, comparing instead the 'conservative' minus 'good' sample (the $\sim$ 1000 QSOs from the 'conservative' sample not classified as 'good') and secondly, constructing two equally sized random samples from the 'conservative' sample. In both instances, we find the same strong correlations as in Figure 2 with similar amplitude variations between the respective correlation matrices as shown in Figure 7.

Finally, our results on line parameter correlations are sensitive to the choice of redshift estimate and any inherent biases in those estimates. We discuss this issue in detail in Appendix A, but, in summary, we find that while different redshift estimates do result in slightly different covariance matrices, our general conclusions are robust and the redshift estimate we have chosen (the BOSS pipeline redshift, $z_{\rm pipeline}$) performs the best at the Lyα profile reconstruction (see the next section).

5 Lyα RECONSTRUCTION

We now use this covariance matrix to reconstruct the Lyα line profile. In this section we outline the reconstruction pipeline, and associated assumptions followed by an example of this approach.

5.1 Reconstruction method

Within our reconstruction approach, we approximate the distribution of the emission line parameters (across the entire data-sample) as a Gaussian[12]. In Figures 3-5, we observe that for these six parameters shown, this approximation is well motivated. Clearly, showing the 18 individual 1D PDFs for each parameters would be uninformative, however, we confirm by eye that this approximation is valid for all model parameters. Note, that in some of these cases, a Gaussian approximation holds only after taking the logarithm, especially for the normalised peak amplitude (normalised by f_{1450}) and the emission line width. Importantly, if anything, these Gaussian approximations have a tendency to slightly overestimate the relative scatter within each parameter (see Figure 5 for example), therefore, this assumption in fact turns out to be a conservative estimate of the true scatter. The Gaussian nature of the scatter in Figures 3-5 and the conservative overestimation of the errors by our Gaussian approach lends confidence that our approach should not significantly underestimate the errors that one might recover from a fully Bayesian approach.

In order to perform the reconstruction, we assume that the QSO can be fit following the same procedure outlined in Section 3, except now we only fit the QSO red-ward of 1275 Å. Our choice of $\lambda > 1275$ Å is conservatively selected to be both close enough to Lyα as possible, but with minimal to no contamination from emission line wings namely Lyα, N V or Si II). Furthermore, given that this approach is best suited for recovering the intrinsic Lyα profile from a Lyα obscured or high-z QSO, it is best to be sufficiently far from any possible contamination of the Lyα line region.

We can then define the N dimensional parameter space (i.e. our 18 emission line parameters outlined previously) as an N dimensional likelihood distribution given by,

$$\mathcal{L} = \frac{1}{(2\pi)^{N/2}|\boldsymbol{\Sigma}|}\exp\left[\frac{1}{2}(\boldsymbol{x}-\boldsymbol{\mu})^{\rm T}\boldsymbol{\Sigma}^{-1}(\boldsymbol{x}-\boldsymbol{\mu})\right]. \quad (7)$$

Here, $\boldsymbol{\Sigma}$ is the recovered QSO covariance matrix (Section 4.1), $\boldsymbol{\mu}$ is the data vector of the means obtained from the full QSO sample for each of the individual line profile parameters and $\boldsymbol{x}$ is the data vector measured from our MCMC fitting algorithm for the individual obscured QSO spectrum.

After the Lyα obscured or high-z QSO has been fit following our fitting procedure, the recovered best-fit values for the unobscured emission line parameters are folded into

[12] The choice in adopting a Gaussian covariance matrix is driven by the large computational burden required to perform a full end-to-end Bayesian approach folding in all modelling uncertainties.

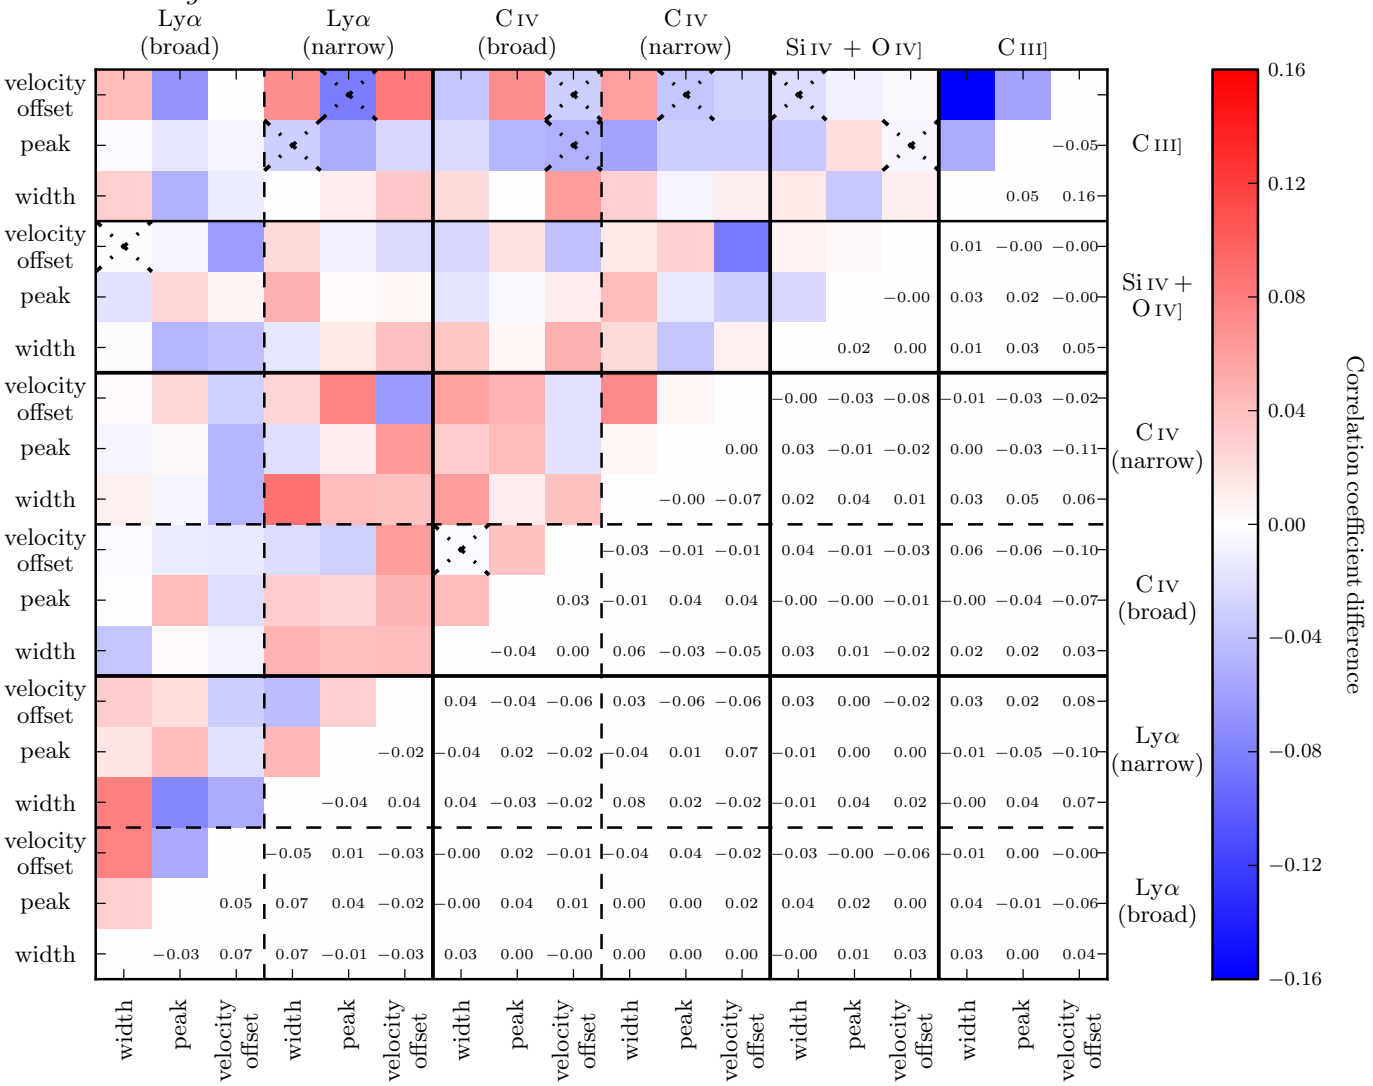

Figure 7. The amplitude of the difference of the correlation coefficients for the 18 × 18 covariance matrix between the 'good' (1673) and 'conservative' (2653) QSO datasets. Positive (red) differences indicate where the correlation in the 'good' sample is greater than the 'conservative' and negative (blue) indicates when the 'conservative' sample is strongest. Note this does not differentiate between a positive or anti-correlation nor the strength of the original correlation, just a strengthening/weakening of the respective correlation. Dot-dashed squares indicate a swap between a positive and negative correlation (typically indicative of the correlation coefficients in either QSO samples being close to zero). The colour bar has been renormalised relative to Figure 2 to indicate the strength of the difference in amplitude.

Equation 7. That is, we recover the best-fit estimates of the Si IV +O IV], C IV and C III] emission lines and evaluate Equation 7 to collapse the 18 dimensional likelihood function into a simple, six dimensional likelihood function describing the six unknown Lyα emission line parameters (two Gaussian components each defined by three parameters). The maximum likelihood of this six dimensional function then describes the best-fit reconstructed profile for the Lyα emission line, while the full six dimensional matrix contains the correlated uncertainty.

5.2 Reconstruction example

We now present an example to highlight the performance of our approach. In order to do this, we choose the same QSO we showed in Figure 1, fitting for $\lambda > 1275$ Å and recovering the six dimensional estimate for the Lyα line profile. In restricting our fitting algorithm to $\lambda > 1275$ Å we are not accessing all the information that was used by the full fit to estimate the QSO continuum. While, this does not affect the recovery of the Lyα peak profile itself, the Lyα profile plus continuum could be affected. In Appendix D we test this assumption, finding that the QSO continuum parameters can be recovered equivalently from these two approaches, with a small amount of scatter in the QSO spectral index.

Before providing the full reconstructed profile, we first recover the individual marginalised 1D PDFs for each of the six Lyα emission line parameters to better visualise the relative size of the errors. In order to obtain the recovered 1D PDFs for the Lyα profile parameters, we marginalise the six dimensional likelihood function over the remaining five Lyα parameters. In Figure 8 we present these recovered 1D marginalised PDFs, showing in the top row the Lyα broad line component and in the bottom row, the nar-

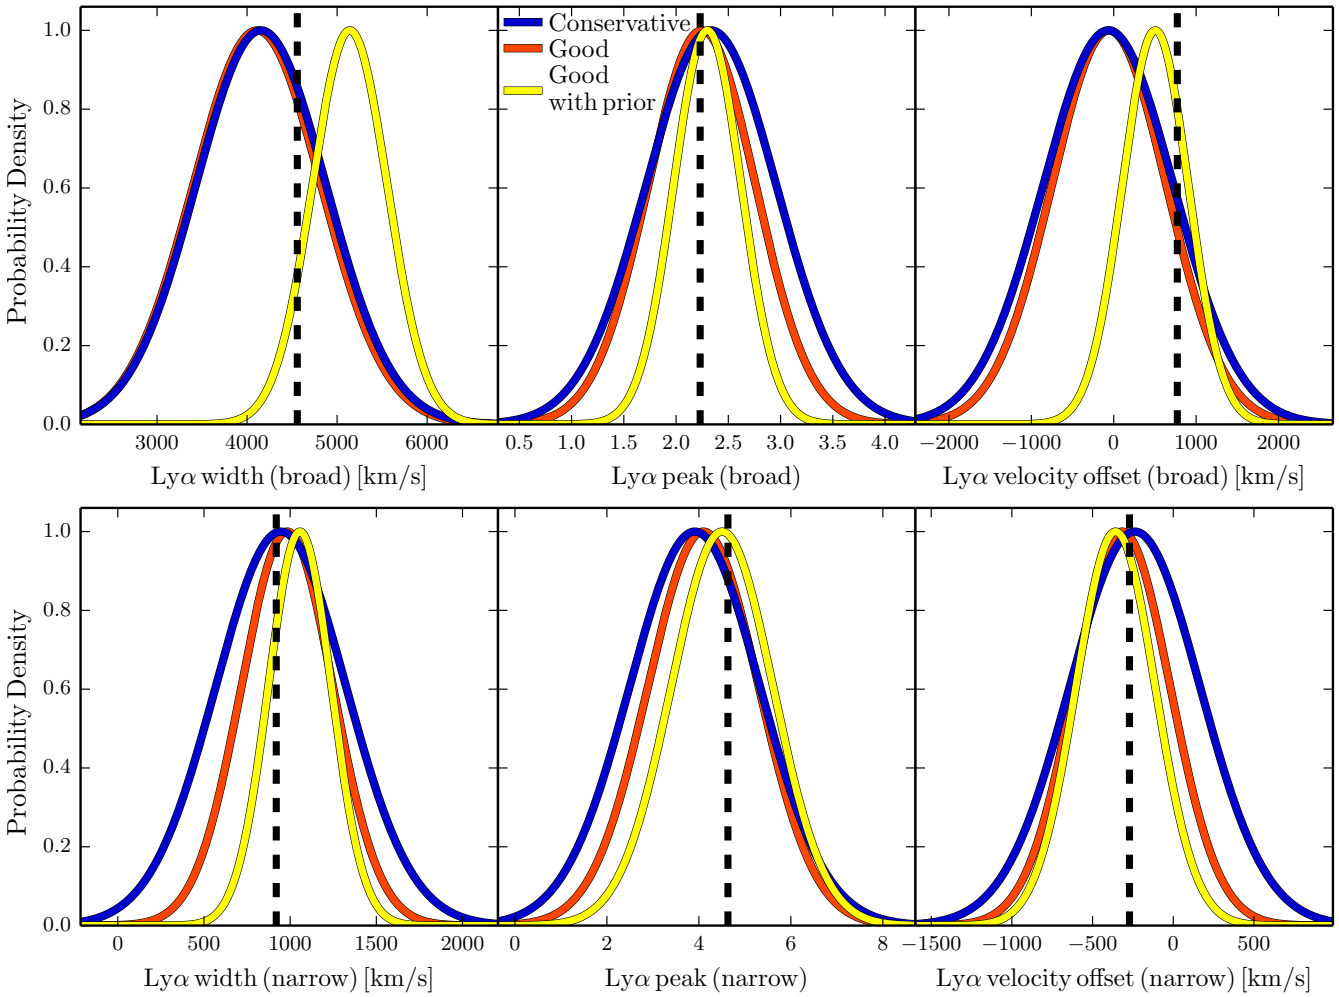

Figure 8. The recovered 1D marginalised PDFs for each of the six reconstructed Lyα emission line parameters obtained after applying the reconstruction method outlined in Section 5.1 to the example QSO in Figure 1. The vertical dashed lines correspond to the MCMC maximum likelihood fit to the full QSO spectrum, whereas the blue and red curves correspond to the recovered 1D PDFs obtained from using the covariance matrix constructed from the 'conservative' and 'good' QSO samples respectively. The yellow curve corresponds to reconstructing the Lyα profile parameters using the 'good' sample while in addition applying a prior on the QSO flux within the range $1230 < \lambda < 1275$Å (see Section 5.3 for further details). That is, we enforce our reconstructed Lyα emission line profile to fit the observed spectrum within this region. Note that the peak amplitude is normalised by f_{1450} and is therefore dimensionless. Importantly, we are interested in the joint probability, i.e. the full Lyα line profile. As shown below, the reconstruction performs considerably better than can be inferred from these marginalised PDFs.

row line component. The vertical black dashed lines represent the recovered values from fitting the full QSO in Figure 1, and the coloured curves represent the recovered 1D marginalised PDFs for each Lyα parameter for the 'good' sample (red) and 'conservative' sample (blue). It is clear that both QSO samples recover almost identical best-fit values for each Lyα parameter, highlighting that both samples are characterised by the same correlations within the covariance matrix. Furthermore, the choice in constructing a 'good' sample is more evident here as consistently the 'good' sample provides marginally narrower constraints. For the most part, the six reconstructed Lyα parameters are recovered within the 68 percentile limits of the original fit to the QSO, with the exception being the velocity offset of the broad Lyα component which is only slightly beyond this limit. Otherwise, the five remaining Lyα parameters are effectively centred around the expected value from the full fit to the same QSO (Figure 1) which should enable a relatively robust recovery of the full Lyα line profile.

In Figure 9 we provide the full reconstructed Lyα line profile. In the top left panel we show the reconstructed Lyα line profile obtained from the 'good' QSO sample, whereas in the top right panel is the reconstructed profile from the 'conservative QSO sample. In all figures, we present 100 reconstructed Lyα profiles denoted by the thin grey curves which are randomly drawn from the full posterior distribution. This small subset of reconstructed profiles highlight the relative scale of the variations in the total Lyα line profile peak height, width and position. The black curve is the raw data from the observed QSO whereas the white curve is the original fit to the full QSO spectrum as shown in Figure 1.

In both, we find the total shape of the reconstructed Lyα line profile to match extremely well with the original full fit to the same QSO, highlighting the strength and utility

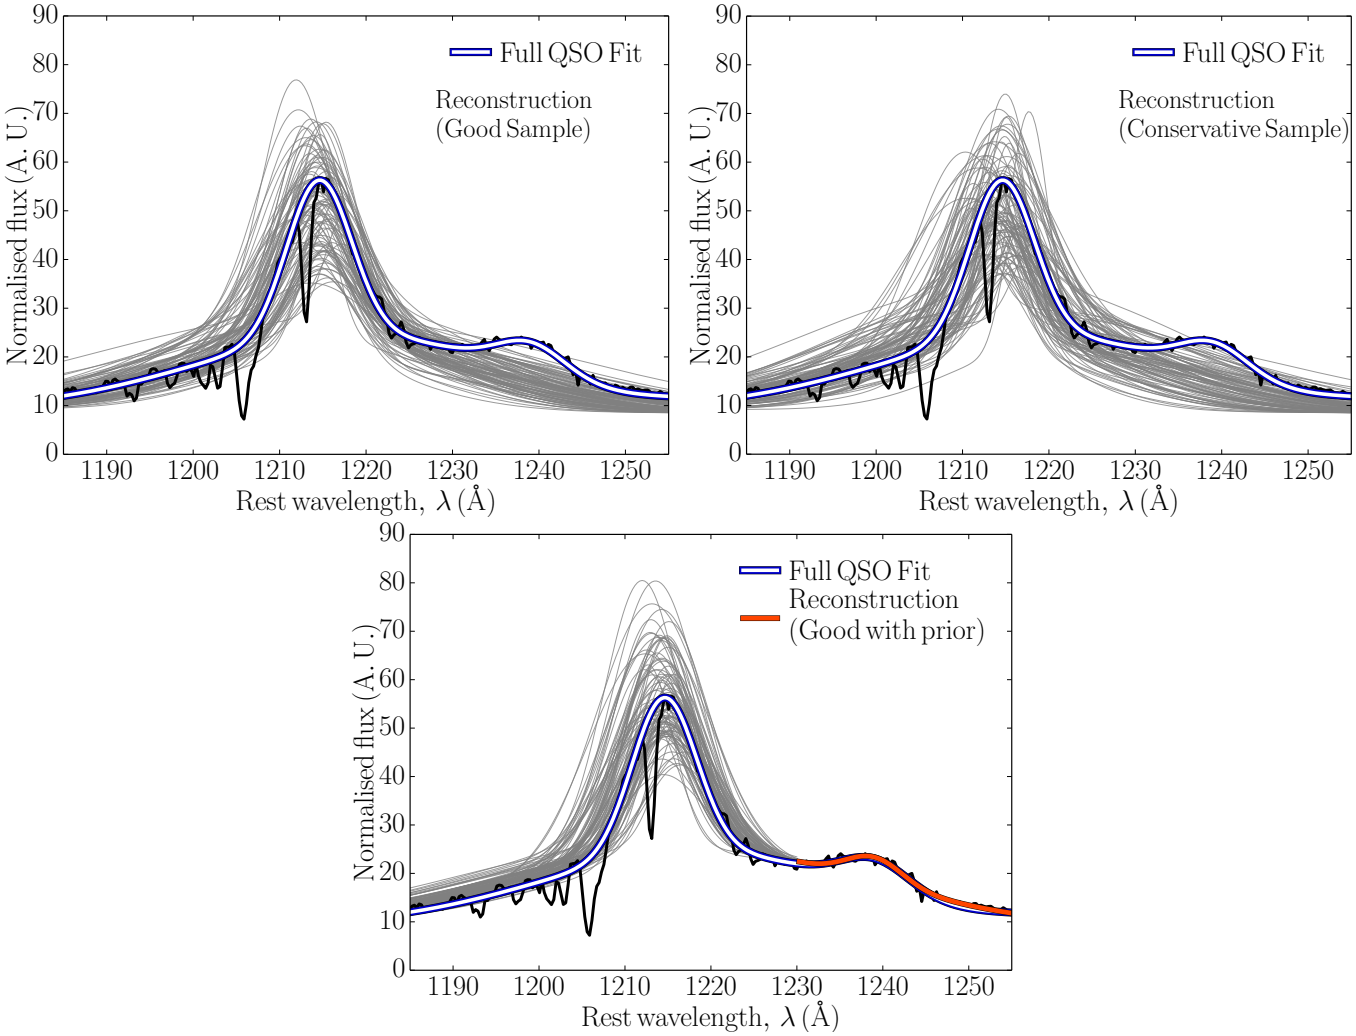

Figure 9. A zoom in of the recovered Lyα emission line profile from our reconstruction procedure. The thin grey curves denote 100 Lyα line profiles extracted from the reconstructed six-dimensional Lyα likelihood function. These curves are randomly selected to represent the full posterior distribution for the reconstructed Lyα profiles, highlighting the relative scale of the errors. The white curves represent the original full MCMC fit of the same QSO (see Figure 1). Black curves are the observed flux of the original QSO spectrum. Note, in these figures, we provide only the emission line component of the full fit (i.e. absorption features are identified and fit, but not shown in the figure). *Top left:* The reconstructed profile from the 'good' sample. *Top right:* The reconstructed profile from the 'conservative' sample. Note that in both the *top left* and *top right* panels the N v emission feature is not fit in the reconstruction procedure. *Bottom:* The reconstructed profile from the 'good' sample utilising the flux prior applied within the range $1230 < \lambda < 1275$Å (see Section 5.3 for further details).

of this covariance matrix reconstruction method. Both the 'good' and 'conservative' samples recover almost the identical reconstructed profile, though the 'conservative' QSO sample provides slightly broader errors. Note that for both these QSO samples, the reconstructed Lyα line profile is systematically below the original fit to the same QSO, indicated by the original QSO fit (white curve) being above the highest density of line profiles[13]. However, this systematic offset is only minor (less than 10 per cent in the normalised flux), well within the errors of the reconstruction. Referring back to Figure 8, we can see that this underestimation appears due to the narrow component of the Lyα peak amplitude (bottom, central panel).

[13] This is more prevalent when averaging over the full distribution of reconstructed Lyα profiles.

5.3 Improving the reconstruction with priors

We presented in Figure 9 our best-fit reconstructed Lyα profiles to a representative QSO drawn from our full sample. However, note that in this reconstruction method we have only used information from the QSO spectrum above $\lambda > 1275$Å. In doing this, for the case of our example QSO we found our maximum likelihood estimates to slightly underestimate (by less than 10 per cent) the original Lyα line profile. At the same time, the total 68 per cent marginalised likelihoods of the reconstructed profiles are relatively broad. Motivated by this, we investigate whether we can provide an additional prior on the Lyα reconstruction profile to further improve the robustness of the recovered profile and to reduce the relative scatter.

In the top panels of Figure 9, redward of Lyα we see that the reconstructed Lyα line profile drops well below the

observed flux of the original QSO. Given we are only reconstructing the Lyα line, this is to be expected as we are not recovering or fitting the N v emission line. However, in both Lyα obscured and high-z QSOs, the N v emission line should be relatively unobscured, therefore, the observed flux within this region could be used as a relative prior on the overall Lyα flux amplitude.

We therefore include a flux prior[14] into our Lyα reconstruction by performing the following steps:

- As before, we fit the QSO at $\lambda > 1275$Å, recovering the QSO continuum and all emission line profiles necessary for our covariance matrix approach.
- Using these estimates, we collapse the 18-dimensional covariance matrix into a six dimensional estimate of the intrinsic Lyα emission line profile.
- We then jointly MCMC sample the observed QSO spectrum within the range $1230 < \lambda < 1275$Å. We fit the N v and Si II lines at the same time sampling from our six dimensional reconstructed Lyα likelihood function obtained from the $\lambda > 1275$Å fit. Fitting to the observed QSO flux, and using the observed noise in the spectrum, we obtain a maximum likelihood for the reconstructed profile. In other words, we require the reconstructed Lyα line profiles to fit the observed spectrum over the range $1230 < \lambda < 1275$Å.

Implementing this prior on the observed flux closer to $\lambda = 1230$Å[15] accesses additional information on the Lyα profile that was not available through our original $\lambda > 1275$Å reconstruction method. Near $\lambda = 1230$Å, there should be a contribution from the Lyα broad line component, which is somewhat degenerate with the N v line (as can be seen in Figure 1). By simultaneously fitting the N v line and the Lyα likelihood function, we use this additional information to place a prior on the Lyα broad line component, which should then reduce the overall scatter in the six dimensional Lyα likelihood function.

Referring back to Figure 8 we provide an example of this prior applied to the same QSO fit and reconstructed previously. In Figure 8, the yellow curves represent the recovered 1D marginalised PDFs for each of the six Lyα emission line parameters. Immediately, it is clear that the application of this prior further reduces the relative error in the recovered PDFs. Furthermore, these PDFs remain centred on the originally recovered values, highlighting we have not biased our reconstruction method.

In the bottom panel of Figure 9 we present the full reconstructed Lyα profile after the addition of this prior. The red curve indicates the fit to the QSO within $1230 < \lambda < 1275$Å, which we have used as our prior to improve the reconstruction of the Lyα line profile. At $\lambda < 1230$Å, we then have the same 100 thin grey curves representing the full posterior distribution of the reconstructed Lyα likelihood profiles. By applying this additional prior on the total observed flux, we reduce the overall scatter in the reconstructed profile. The maximum likelihood profiles (thin grey curves) now provide a more robust match to the observed QSO.

5.4 Statistical performance of the reconstruction pipeline

Thus far we have only applied our reconstruction pipeline to recover the Lyα line profile from a single, test example selectively drawn from our 'good' sample of 1673 QSOs. In order to statistically characterise the performance of the Lyα reconstruction pipeline across the QSO sample, in Figure 10 we present 2D scatter plots of the maximum a posteriori (MAP) estimates[16] of the reconstructed Lyα emission line parameters compared to the originally obtained values from the full MCMC fit to the QSO (without masking out the Lyα line). For each of the six Lyα emission line parameters, we highlight the 68 and 95 percentiles of the joint marginalised likelihoods for the distributions by the green solid and dashed contours, respectively. Additionally, the red dashed curve demarcates the one-to-one line, along which all QSOs would sit if the reconstruction profile worked idealistically.

Across the six Lyα panels, we find strong agreement ($\rho > 0.7$) amongst half of our Lyα line parameters, those being the peak amplitudes of both the Lyα broad and narrow components and the velocity offset for the Lyα narrow line component. For the remainder of the parameters we find moderate to weaker recovery of the original line parameters. However, note that in order to keep this figure as clear as possible we are only providing the MAP estimates. Within Figures 8 and 9 we found that the relative scatter on the reconstructed Lyα line profile parameters was notable, and therefore the majority of the reconstructed parameters are within the 68 per cent marginalised errors.

The reconstructed Lyα profile parameters highlighted here reflect the correlations recovered from the covariance matrix in Figure 2. We found strong correlations in the peak amplitudes of the Lyα profiles, and the narrow component velocity offset. The lack of a strong correlation for the width of the Lyα line in our covariance matrix, translates to weaker recovery of these parameters. In principle, these weaker correlations could be further strengthened by adding an appropriate prior on the line widths motivated by the statistical distributions recovered from the full sample, or other line properties such as correlations between the equivalent widths.

In Figure 10, there is also slight evidence for a bias in the reconstructed parameters, as highlighted by the orientation of the green contours (68 and 95 percentiles of the reconstructed parameter distributions) relative to the reference one-to-one line. However, this could artificially arise as the increase/decrease in any one of these Lyα line parameters could be compensated for by respective changes in others (i.e. model degeneracies), whereas the full six dimensional

[14] Note that this choice in phrasing does not imply a 'prior' in the Bayesian statistical sense, rather it is used to highlight that we are including additional information into our reconstruction procedure compared to the reconstruction method discussed in Section 5.1.

[15] This choice of 1230Å is purely arbitrary, and is chosen based on the assumption that a damped absorption signal would not extend this far redward of Lyα. However, this choice is flexible, and can be adjusted on a case-by-case basis if evidence for stronger attenuation beyond 1230Å is present.

[16] Note that throughout this work, the recovered MAP estimates from the full posterior distribution do not differ significantly from the peak of the associated marginalised PDFs.

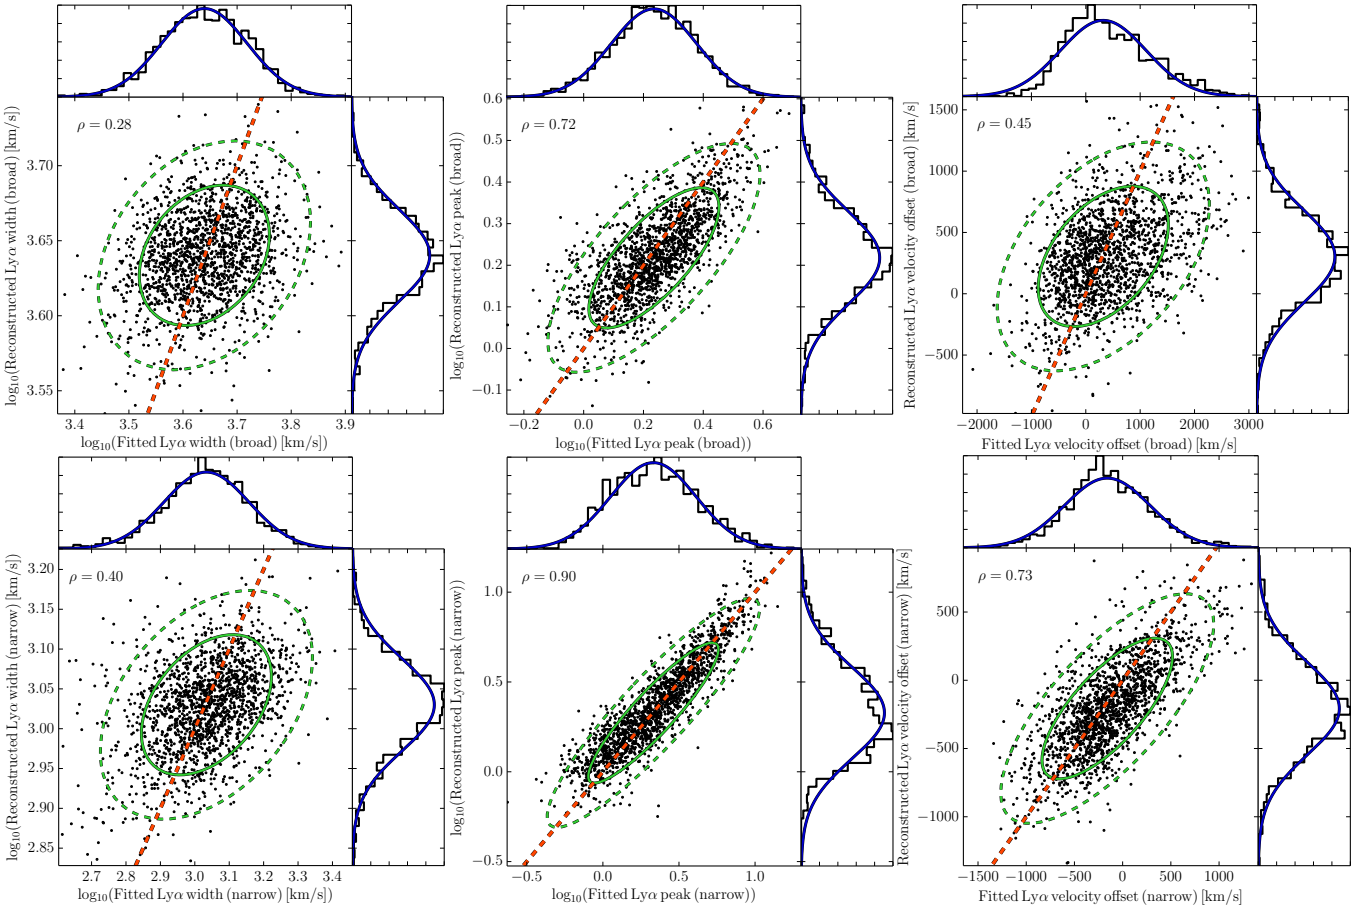
ρ = 0.28
log10(Reconstructed Lyα width (broad) [km/s])
log10(Fitted Lyα width (broad) [km/s])
ρ = 0.72
log10(Reconstructed Lyα peak (broad))
log10(Fitted Lyα peak (broad))
ρ = 0.45
Reconstructed Lyα velocity offset (broad) [km/s]
Fitted Lyα velocity offset (broad) [km/s]
ρ = 0.40
log10(Reconstructed Lyα width (narrow) [km/s])
log10(Fitted Lyα width (narrow) [km/s])
ρ = 0.90
log10(Reconstructed Lyα peak (narrow))
log10(Fitted Lyα peak (narrow))
ρ = 0.73
Reconstructed Lyα velocity offset (narrow) [km/s]
Fitted Lyα velocity offset (narrow) [km/s]

Figure 10. A visual characterisation of the performance of the Lyα emission line reconstruction pipeline at recovering each of the six Lyα line parameters. We provide the 2D scatter between the full MCMC fits to the QSOs, compared to the reconstructed values from a fit to the same QSO masking the Lyα region (flux masked at $\lambda < 1275$Å) and applying our additional prior on the QSO flux near $\lambda = 1230$Å. Green solid and dashed contours enclose the 68 and 95 per cent scatter of the reconstructed Lyα parameters relative to their expected (true) value, while the red dashed curves correspond to the one-to-one line on which all points would lie if the reconstruction procedure worked perfectly. Histograms (black curves) correspond to the 1D PDFs of the Lyα parameters from the full MCMC fitting (top) and the recovered estimate from the reconstruction pipeline fit to the Lyα masked QSO (right). Blue curves represent the associated Gaussian distribution with equivalent scatter. Note that the peak amplitude is normalised by f_{1450} and is therefore dimensionless.

Lyα likelihood function takes these model degeneracies into account when estimating the full reconstructed Lyα profile.

In order to better illustrate the full reconstruction of the joint Lyα parameter likelihoods, in Figure 11 we show the information from the six individual Lyα line parameters as a single, total measured Lyα line flux at two arbitrarily defined locations blue (1205Å) and redward (1220Å) of the Lyα line centre. We compare the total reconstructed Lyα line flux against the measured Lyα line flux from our full fit to the same QSO, providing the reference one-to-one relation as the red dashed curve, and the grey shaded region encompasses the region in which the reconstructed Lyα line flux is within 15 per cent of the measured flux. Immediately obvious from this figure is that there are no apparent biases in the reconstruction process, i.e. we neither systematically over nor underestimate the reconstructed Lyα line profile. This figure highlights the strength of the Lyα reconstruction process. Redward of Lyα (at 1220Å), closer to our flux prior at 1230Å, we find that $\sim$ 90 per cent of all reconstructed Lyα line profiles have a recovered flux within 15 per cent. As one would expect, the scatter increases blueward of Lyα (at 1205Å), however we still find the reconstructed Lyα line flux to be within 15 per cent for $\sim$ 85 per cent of our sampled QSOs. This highlights statistically, that the reconstruction process performs an excellent job of recovering the full Lyα line profile.

6 POTENTIAL APPLICATIONS

In this work, we have developed an MCMC fitting algorithm for the sole purpose of characterising the QSO continuum and the emission line profiles within the range 1180Å $< \lambda <$ 2300Å. Our goal was the construction of a covariance matrix to reconstruct the Lyα line profile. However, due to the flexibility of the MCMC approach, many other applications could benefit from such a pipeline.

Firstly, we were only interested in correlations amongst the strongest, high ionisation emission line parameters for our covariance matrix. However, various properties of QSOs can be extracted from accurate recovery of the line widths and ratios. For example, the QSO metallicity has been estimated from measuring the N V/C IV, N V/He II and the Si IV +O IV]/C IV line ratios in samples of QSOs from Na-

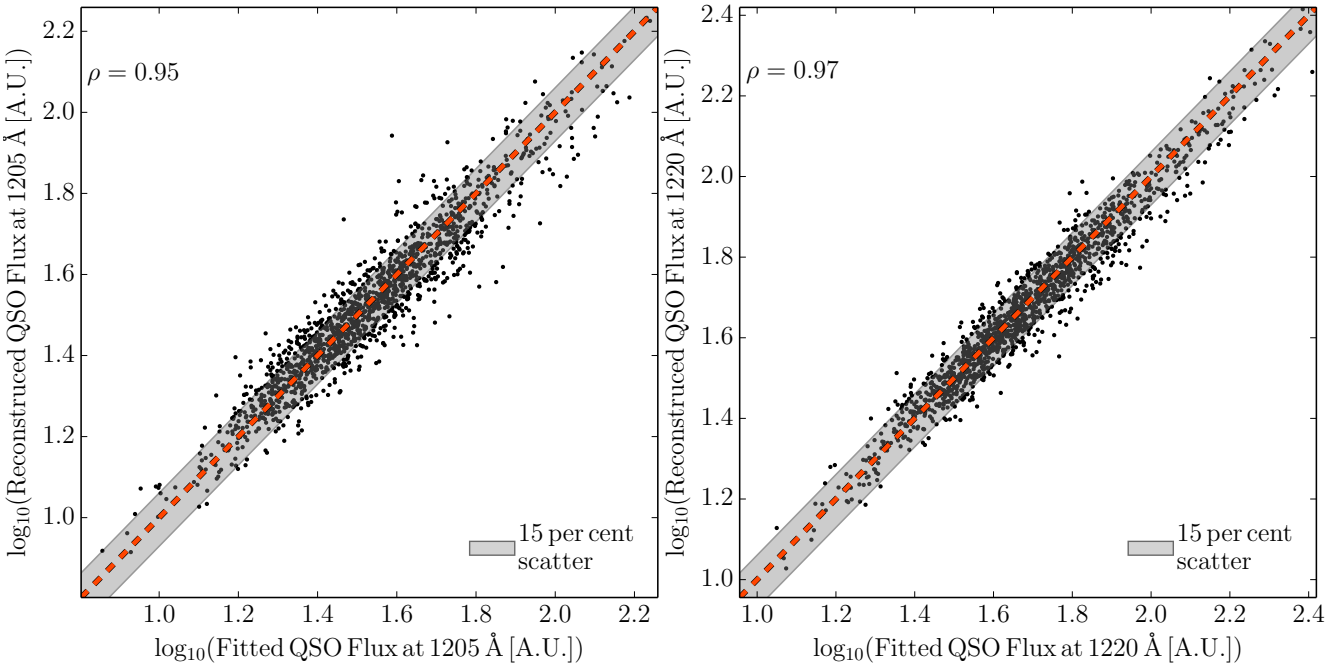
ρ = 0.95
15 per cent scatter
log10(Reconstruced QSO Flux at 1205 Å [A.U.])
log10(Fitted QSO Flux at 1205 Å [A.U.])
ρ = 0.97
15 per cent scatter
log10(Reconstruced QSO Flux at 1220 Å [A.U.])
log10(Fitted QSO Flux at 1220 Å [A.U.])

Figure 11. A comparison of the maximum likelihood reconstructed Lyα line flux (including the flux prior; Section 5.3) to the actual Lyα line flux obtained from the full fit at a specific wavelength for each QSO within our 'good' sample (1673 QSOs; we define 1 A. U. = 10^{-17} erg cm^{-2} s^{-1} Å^{-1}). *Left panel:* blueward of Lyα (λ = 1205Å). *Right panel:* redward (λ = 1220Å). The red dashed curve corresponds to the one-to-one relation, and the grey shaded region encompasses 15 per cent scatter in the reconstructed Lyα flux relative to the actual measured Lyα line flux. At λ = 1220Å we find the reconstructed Lyα line flux to be within 15 per cent of the actual Lyα line flux in $\sim$ 90 per cent of all QSOs, decreasing to $\sim$ 85 per cent at λ = 1205Å.

gao et al. (2006b) ($2 < z < 4.5$) and Juarez et al. (2009) ($4 < z < 6.4$). Using the same line ratios, we could recover estimates for the metallicities for all QSOs within our measured sample.

Several other emission line ratios (e.g. the R23 parameter, [O III]/[O II], C IV/He II, N-based ratios) can additionally be used as proxies for QSO metallicity (e.g Nagao et al. 2006a; Matsuoka et al. 2009; Batra & Baldwin 2014). By extending our QSO MCMC framework, many other emission lines can be obtained and characterised to improve the QSO metallicity estimates. Crucially, this MCMC approach, would enable a large dataset of QSOs to be rapidly explored. The extension to measuring the metallicities of the QSOs would enable quantities such as the SMBH mass to be recovered through existing correlations between the emission line FWHMs.

In Section 3.3 we have outlined our method to identify and fit absorption features (see e.g. Zhu & Ménard 2013 for a more robust approach). While in this work these have been considered contaminants, the prevalence and measurement of these lines could be used to infer properties of the metallicities within intervening absorbers (e.g Ryan-Weber et al. 2006; Becker et al. 2009; D'Odorico et al. 2013) and to reveal the presence of massive outflows of ionised gas from their nuclei (e.g. Crenshaw et al. 2003, and reference therein). For example, we have thrown out any QSOs with strong intervening absorption from systems such as DLAs, however, analysing these sources with our MCMC fitting algorithm could yield measurements on the internal properties of these absorption systems.

Assuming the input quasars to our MCMC are representative of the quasar population as a whole, and further that the UV spectral properties of quasars do not evolve with redshift, our model can be used to predict the intrinsic distribution of quasar spectra at redshifts beyond those from which the model was calibrated ($z \sim 2.5$). A similar procedure has been often employed to characterise the colour selection efficiency of quasar surveys (e.g Fan 1999), although the correlations obtained through our MCMC approach provide more detailed reconstruction of quasar emission features and hence more reliable colour models. Particularly since the Lyα line plays such a key role in the selection of high-z quasars, our model could be used to identify quasars missing from current surveys due to selection effects and provide more robust statistics for high-z quasar luminosity functions.

In addition, by recovering an estimate for the intrinsic Lyα emission line profile one can investigate the QSO proximity effect. This approach requires an estimate of the intrinsic QSO luminosity, coupled with the modelling of the Lyα forest. Within the sphere of influence of the QSO, the photoionisation background is higher than the mean background permeating the IGM. Modelling the transition between the mean IGM background and the drop-off in QSO luminosity has been used by several authors to recover estimates of the photoionisation background in the IGM (e.g Bolton et al. 2005; Bolton & Haehnelt 2007b; Calverley et al. 2011). At the redshifts where these studies have been performed (e.g. $2 < z < 6$), the Lyα line profile is not obscured or attenuated by a neutral IGM. Therefore, one could in principle push the flux prior we used in this work much closer to the Lyα line centre, to substantially reduce the errors on the intrinsic Lyα emission line profile blueward of Lyα.

At $z > 6$, for an increasingly neutral IGM, the intrinsic

Lyα emission line can become increasingly attenuated by the Gunn-Peterson IGM damping wing. While existing methods have been developed to access information on the red-side of Lyα (e.g. Kramer & Haiman 2009), for the $z = 7.1$ QSO ULASJ1120+0641 (Mortlock et al. 2011), evidence suggests that the IGM damping wing imprint extends further redward (e.g Mortlock et al. 2011; Bolton et al. 2011), limiting the effectiveness of these approaches. The approach developed in this work should be unaffected by this as here we do not fit the Lyα line, our reconstruction method therefore is perfectly suited for exploring the potential imprint of the IGM damping wing on the $z = 7.1$ QSO (e.g. Greig et al. 2016), along with other future $z > 6$ QSOs.

7 CONCLUSION

Characterising the continuum and emission lines properties of QSOs provides a wealth of information on the internal properties of the AGN, such as the mass of the SMBH, the QSO metallicity, star formation rates of the host galaxy, nuclear outflows and winds etc. Furthermore, the intrinsic Lyα line shape can be used to probe properties of the IGM, such as the mean photoionisation background and the abundance of neutral hydrogen in the IGM at $z > 6$.

Motivated by correlations amongst QSO emission lines (e.g. Boroson & Green 1992; Sulentic et al. 2000; Shen et al. 2011), in this work, we developed a new reconstruction method to recover the intrinsic Lyα profile. This method is based on the construction of a covariance matrix built from a large sample of moderate-z ($2.0 < z < 2.5$), high S/N (S/N > 15) QSOs from the BOSS observational programme. We use this moderate-z sample to characterise the intrinsic Lyα line profile, where it should be relatively unaffected by intervening neutral hydrogen in the IGM. In order to characterise each QSO within our sample we developed an MCMC fitting algorithm to jointly fit the QSO continuum, the emission lines and any absorption features that could contaminate or bias the fitting of the QSO. We modelled the QSO continuum as a single two parameter power-law ($\propto \lambda^{\alpha_\lambda}$), and each emission line is modelled as a Gaussian defined by three parameters, its width, peak amplitude and velocity offset from systemic.

We constructed our covariance matrix from a refined sample of 1673 QSOs, using the high ionisation emission lines Lyα, Si IV +O IV], C IV and C III]. Owing to the flexibility of the MCMC framework, we explored various combinations of single and double component Gaussians to characterise these lines. For Lyα and C IV we settled on a double component Gaussian, to model the presence of a broad and narrow component. For the remaining two lines, we considered a single component. This resulted in an 18×18 covariance matrix, which we used to investigate new and existing correlations amongst the line profiles. We identified several strong trends from our covariance matrix, most notably the strong positive correlation between the peak amplitudes of the Lyα and C IV narrow components ($\rho = 0.8$) and the strong anti-correlation between the Lyα narrow component peak amplitude and the width of the C III] line ($\rho = -0.74$). These two were the strongest examples of a consistent trend of a positive correlation in the peak amplitudes across all the emission line species and the anti-correlation between the peak amplitudes and line widths.

Using this covariance matrix, we constructed an N-dimensional Gaussian likelihood function from which we are able to recover our reconstructed Lyα line profile. The reconstruction method works as follows:

- Fit a QSO with our MCMC pipeline within the range $1275\text{Å} < \lambda < 2300\text{Å}$, recovering the parameters defining the QSO continuum, and the Si IV +O IV], C IV and C III] lines.
- Obtain a six dimensional estimate of the reconstructed Lyα line profile (modelled as a two component Gaussian) which provides the best-fit profile and correlated uncertainties, by evaluating the N-dimensional likelihood function describing our full covariance matrix including a prior on the reconstructed Lyα line using the observed QSO flux within the range $1230 < \lambda < 1275\text{Å}$.

To visually demonstrate the performance of this reconstruction method, we applied it to a randomly selected QSO from our full data set. Finally, we quantitatively assessed its performance by applying it to the full QSO sample, and compared the reconstructed Lyα profile parameters to those recovered from the original full MCMC fit of the same QSO. We found that estimates for both the Lyα peak amplitudes are recovered strongly, as is the velocity offset of the narrow line component. For both of the line widths and the broad component velocity offset we find moderate agreement. We additionally explored the total reconstructed Lyα flux (rather than individual parameters) relative to the original full MCMC fit at two distinct wavelengths blueward (1205Å) and redward (1220Å) of Lyα. Our reconstruction method recovered the Lyα line flux to within 15 per cent of the measured flux at 1205Å (1220Å) ~ 85 (~ 90) per cent of the time.

There are several potential applications for both the MCMC fitting method and the Lyα reconstruction pipeline. The MCMC fitting could be easily modified to measure any emission or absorption feature within a QSO spectrum. With this, many properties of the source QSO could be extracted, for example QSO metallicities. The ability to reconstruct the intrinsic Lyα line profile could have important cosmological consequences such as improving estimates of the IGM photoionisation background or recovering estimates of the IGM neutral fraction.

ACKNOWLEDGMENTS

We thank the anonymous referee for their helpful suggestions. AM and BG acknowledge funding support from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 638809 – AIDA – PI: AM). ZH is supported by NASA grant NNX15AB19G.

Funding for SDSS-III has been provided by the Alfred P. Sloan Foundation, the Participating Institutions, the National Science Foundation, and the U.S. Department of Energy Office of Science. The SDSS-III web site is http://www.sdss3.org/.

SDSS-III is managed by the Astrophysical Research Consortium for the Participating Institutions of the SDSS-III Collaboration including the University of Arizona, the

Brazilian Participation Group, Brookhaven National Laboratory, Carnegie Mellon University, University of Florida, the French Participation Group, the German Participation Group, Harvard University, the Instituto de Astrofisica de Canarias, the Michigan State/Notre Dame/JINA Participation Group, Johns Hopkins University, Lawrence Berkeley National Laboratory, Max Planck Institute for Astrophysics, Max Planck Institute for Extraterrestrial Physics, New Mexico State University, New York University, Ohio State University, Pennsylvania State University, University of Portsmouth, Princeton University, the Spanish Participation Group, University of Tokyo, University of Utah, Vanderbilt University, University of Virginia, University of Washington, and Yale University.

REFERENCES

Akeret J., Seehars S., Amara A., Refregier A., Csillaghy A., 2013, Astron. Comput., 2, 27
Alam S., et al., 2015, ApJS, 219, 12
Baldwin J., Ferland G., Korista K., Verner D., 1995, ApJ, 455, L119
Baldwin J. A., 1977, ApJ, 214, 679
Baldwin J. A., et al., 1996, ApJ, 461, 664
Barkana R., 2002, New Astron., 7, 85
Batra N. D., Baldwin J. A., 2014, MNRAS, 439, 771
Becker G. D., Hewett P. C., Worseck G., Prochaska J. X., 2013, MNRAS, 430, 2067
Becker G. D., Rauch M., Sargent W. L. W., 2009, ApJ, 698, 1010
Blandford R. D., McKee D. F., 1982, ApJ, 255, 419
Bolton A. S., et al., 2012, AJ, 144, 144
Bolton J. S., Haehnelt M. G., 2007a, MNRAS, 374, 493
—, 2007b, MNRAS, 382, 325
Bolton J. S., Haehnelt M. G., Viel M., Springel V., 2005, MNRAS, 357, 1178
Bolton J. S., Haehnelt M. G., Warren S. J., Hewett P. C., Mortlock D. J., Venemans B. P., McMahon R. G., Simpson C., 2011, MNRAS, 416, L70
Boroson T. A., Green R. F., 1992, ApJS, 80, 109
Bosman S. E. I., Becker G. D., 2015, MNRAS, 452, 1105
Bovy J., et al., 2011, ApJ, 729, 141
Brotherton M. S., Tran H. D., Becker R. H., Gregg M. D., Laurent-Muehleisen S. A., White R. L., 2001, ApJ, 546, 775
Calverley A. P., Becker G. D., Haehnelt M. G., Bolton J. S., 2011, MNRAS, 412, 2543
Carilli C. L., et al., 2010, ApJ, 714, 834
Corbin M. R., Boroson T. A., 1996, ApJS, 107, 69
Crenshaw D. M., Kraemer S. B., George I. M., 2003, ARA&A, 41, 117
Dawson K. S., et al., 2013, AJ, 145, 10
Dietrich M., Hamann F., Shields J. C., Constantin A., Vestergaard M., Chaffee F., Foltz C. B., Junkkarinen V. T., 2002, ApJ, 581, 912
Dijkstra M., 2014, PASA, 31, e040
Dodelson S., Schneider M. D., 2013, Phys. Rev. D, 88, 063537
D'Odorico V., et al., 2013, MNRAS, 435, 1198
Eisenstein D. J., et al., 2011, AJ, 142, 72
Fan X., 1999, AJ, 117, 2528
Fan X., et al., 2006, AJ, 132, 117
Fitzpatrick E. L., 1999, PASP, 111, 63
Font-Ribera A., et al., 2013, JCAP, 05, 018
Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, PASP, 125, 306
Francis P. J., Hewett P. C., Foltz C. B., Chaffee F. H., 1992, ApJ, 398, 476
Francis P. J., Hewett P. C., Foltz C. B., Chaffee F. H., Weymann R. J., Morris S. L., 1991, ApJ, 373, 465
Gallerani S., Ferrara A., Fan X., Choudhury R. T., 2008, MNRAS, 386, 359
Goodman J., Weare J., 2010, Commun. Appl. Math. Comput. Sci., 5, 1
Greig B., Mesinger A., Haiman Z., Simcoe R. A., 2016, preprint (arXiv:1606.00441)
Gunn J. E., et al., 2006, AJ, 131, 2332
Harris D. W., et al., 2016, AJ, 151, 155
Hewett P. C., Wild V., 2010, MNRAS, 405, 2302
Hopkins P. F., et al., 2004, AJ, 128, 1112
Juarez Y., Maiolino R., Mujica R., Pedani M., Marinoni S., Nagao T., Marconi A., Oliva E., 2009, A&A, 494, L25
Kaspi S., Smith P. S., Netzer H., Maoz D., Jannuzi B. T., Giveon U., 2000, ApJ, 533, 631
Kirkpatrick J. A., Schlegel D. J., Ross N. P., Myers A. D., Hennawi J. F., Sheldon E. S., Schneider D. P., Weaver B. A., 2011, ApJ, 743, 125
Kramer R. H., Haiman Z., 2009, MNRAS, 400, 1493
Lee K.-G., Spergel D. N., 2011, ApJ, 734, 21
Liddle A. R., 2004, MNRAS, 351, L49
Margala D., Kirkby D., Dawson K., Bailey S., Blanton M., Schneider D. P., 2016, ApJ, 831, 157
Matsuoka K., Nagao T., Maiolino R., Marconi A., Taniguchi Y., 2009, A&A, 503, 721
McDonald P., Eisenstein D. J., 2007, Phys. Rev. D, 76, 63009
McGreer I. D., Mesinger A., D'Odorico V., 2015, MNRAS, 447, 499
McLean I. S., Steidel C. C., Epps H., et al., 2010, Proc. SPIE, 7735, 77351E
McLean I. S., Steidel C. C., Epps H. W., et al., 2012, Proc. SPIE, 8446, 84460J
Mesinger A., 2010, MNRAS, 407, 1328
Mesinger A., Haiman Z., 2004, ApJ, 611, L69
—, 2007, ApJ, 660, 923
Miralda-Escudé J., 1998, ApJ, 501, 15
Mortlock D. J., et al., 2011, Nature, 474, 616
Nagao T., Maiolino R., Marconi A., 2006a, A&A, 459, 85
Nagao T., Marconi A., Maiolino R., 2006b, A&A, 447, 157
Pâris I., et al., 2011, A&A, 530, A50
—, 2016, preprint (arXiv:1608.06483)
Peebles P. J. E., 1993, Principles of Physical Cosmology. Princeton Univ. Press, Princeton, NJ
Peterson B. M., 1993, PASP, 105, 1084
Peterson B. M., et al., 2004, ApJ, 613, 682
Petri A., May M., Haiman Z., 2016, Phys. Rev. D, 94, 063534
Planck Collaboration XIII, 2016, A&A, 594, 13
Planck Collaboration XLVI, 2016, preprint (arXiv:1605.02985)
Planck Collaboration XLVII, 2016, preprint (arXiv:1605.03507)
Rauch M., 1998, ARA&A, 36, 267

Richards G. T., et al., 2003, AJ, 126, 1131
—, 2011, AJ, 141, 167
Ross N. P., et al., 2012, ApJS, 199, 3
Ryan-Weber E. V., Pettini M., Madau P., 2006, MNRAS, 371, L78
Schroeder J., Mesinger A., Haiman Z., 2013, MNRAS, 428, 3058
Schwarz G., 1978, Annals of Statistics, 461, 6
Scott J. E., Kriss G. A., Brotherton M., Green R. F., Hutchings J., Shull J. M., Zheng W., 2004, ApJ, 615, 135
Shang Z., Wills B. J., Wills D., Brotherton M. S., 2007, AJ, 134, 294
Shen Y., Greene J. E., Strauss M. A., Richards G. T., Schneider D. P., 2008, ApJ, 680, 169
Shen Y., Ménard B., 2012, ApJ, 748, 131
Shen Y., et al., 2011, ApJS, 194, 45
Shull J. M., Stevans M., Danforth C. W., 2012, ApJ, 752, 162
Smee S. A., et al., 2013, AJ, 146, 32
Stevans M., Shull J. M., Danforth C. W., Tilton E. M., 2014, ApJ, 794, 75
Sulentic J. W., Marziani P., Dultzin-Hacyan D., 2000, ARA&A, 38, 521
Suzuki N., 2006, ApJS, 163, 110
Suzuki N., Tytler D., Kirkman D., O'Meara J. M., Lubin D., 2005, ApJ, 618, 592
Taylor A., Joachimi B., 2014, MNRAS, 442, 2728
Telfer R. C., Zheng W., Kriss G. A., Davidsen A. F., 2002, ApJ, 565, 773
Vanden Berk D. E., et al., 2001, AJ, 122, 549
Vernet J., et al., 2011, A&A, 536, 105
Vestergaard M., Peterson B. M., 2006, ApJ, 641, 689
Vestergaard M., Wilkes B. J., 2001, ApJS, 134, 1
White M., 2003, The Davis Meeting On Cosmic Inflation. 2003 March 22-25, 18, arXiv:astro-ph/0305474
Wills B. J., Brotherton M. S., Fang D., Steidel C. C., Sargent W. L. W., 1993, ApJ, 415, 563
Wyithe J. S. B., Bolton J. S., 2011, MNRAS, 412, 1926
Wyithe J. S. B., Loeb A., 2004, ApJ, 612, 597
Zhu G., Ménard B., 2013, ApJ, 770, 130

APPENDIX A: SELECTING THE SYSTEMIC REDSHIFT FOR OUR BOSS SAMPLE

The reconstruction method discussed within this work is strongly dependant on having an accurate systemic redshift for the QSO sample. In the absence of knowing the true source redshift, systematic errors in its estimation will filter into the recovered line profile parameters, most notably the velocity offset of the line centre. As a result, this will affect any reported line profile correlations which are necessary for the Lyα line profile reconstruction. Within this appendix, we discuss our attempts to determine the most robust and accurate redshift to be adopted within this work.

The BOSS DR12Q database (Pâris et al. 2016) provides six available redshift estimates; a visually inspected redshift, $z_{\rm VI}$, the BOSS pipeline redshift based on a decomposed eigenvalue training set, $z_{\rm pipeline}$ (see e.g. Bolton et al. 2012), an automated PCA based approach, $z_{\rm PCA}$, and three other estimates determined from the location of the peak emission for the individual line profiles, $z_{\rm C\,IV}$, $z_{\rm C\,III]}$ and $z_{\rm Mg\,II}$. Using the earlier DR9Q database, Font-Ribera et al. (2013) performed a detailed analysis on each of these redshift estimates. These authors found $z_{\rm PCA}$ to be the optimal redshift estimator, returning both the smallest dispersion and bias amongst these six available measurements.

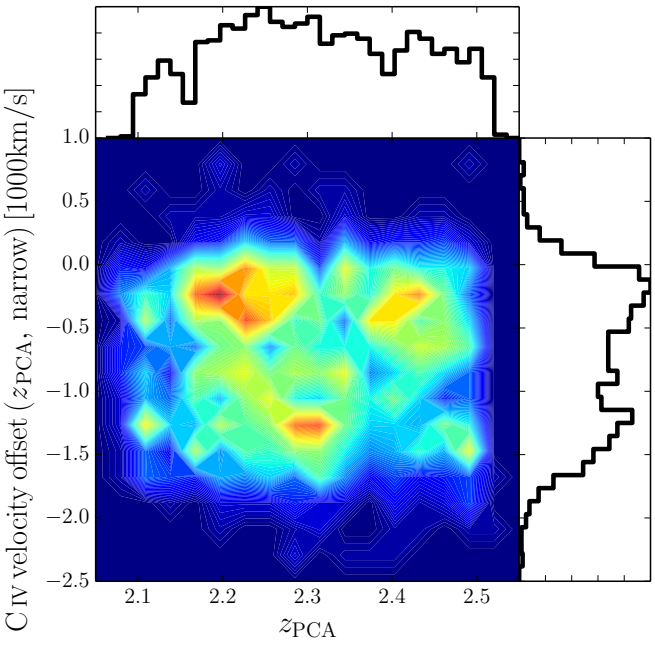

Figure A1. Distribution of the velocity offset from the C IV narrow component against the estimated PCA redshift, $z_{\rm PCA}$ for our 'good' sample of 1673 QSOs. In the central panel is a two-dimensional histogram normalised to unity, with red denoting the highest density of QSOs. In the top (right) panels are the one-dimensional histograms for the distribution of $z_{\rm PCA}$ (C IV narrow line component velocity offset).

However, in adopting $z_{\rm PCA}$ as the optimal redshift estimator we find a bias within our selected sample of QSOs. In Figure A1, we show the velocity offset of the narrow component of the C IV line (measured with $z_{\rm PCA}$) as a function of the estimated redshift from the PCA approach using our 'good' sample of 1673 QSOs. In the right most panel, we show the one-dimensional histogram for the narrow component of C IV. Immediately obvious from this figure, is the non-Gaussian (bi-modal) nature of the distribution. Contrast this with the equivalent figure of the narrow component of C IV for the $z_{\rm pipeline}$ redshift estimator shown in Figure A2, which shows no obvious departure from Gaussianity. We have additionally tested this for the other redshift estimators (e.g. $z_{\rm VI}$ and $z_{\rm Mg\,II}$) and found no evidence for the same bi-modal distribution as shown by $z_{\rm PCA}$.

In the central panel of Figure A1, we show the two-dimensional histogram, with red signifying the highest density of QSOs. Presenting the data in this manner immediately highlights the source of error. The two peaks in the C IV velocity offset distribution are driven by two separate, discrete regions within the C IV–$z_{\rm PCA}$ parameter space: the primary peak in the C IV velocity offset at ~ -250 km/s (for all redshifts except the 'hole' at $z \sim 2.3$) and the secondary peak at ~ -1250 km/s (at $z \sim 2.3$). This discrete 'overdensity' of QSOs at $z \sim 2.3$ is most likely an unphysical artefact. The likely cause of this discrete transition near $z \sim 2.3$ is the redshifting of the Mg II line between strong

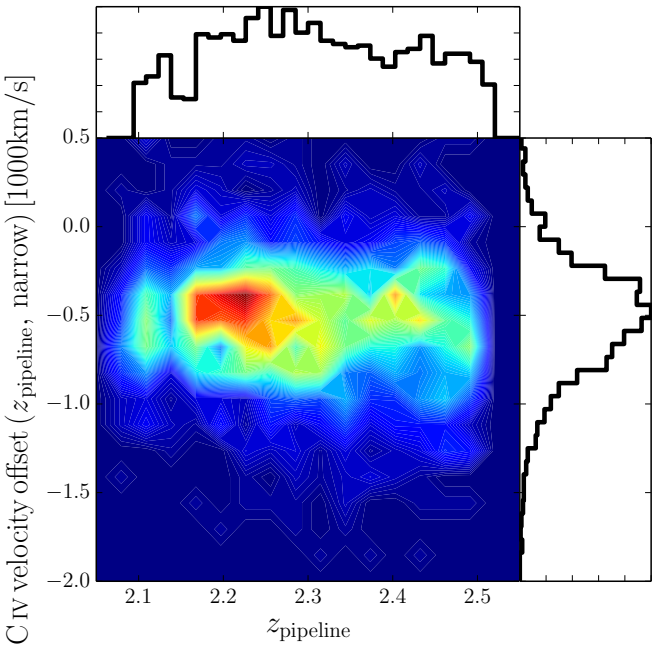

Figure A2. Same as Figure A1, except showing the C IV narrow component measured from the $z_{\rm pipeline}$ redshift estimator.

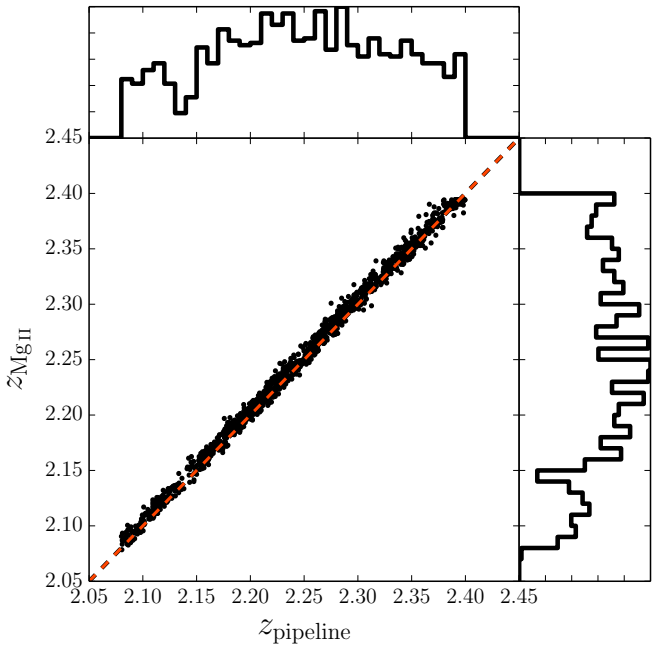

Figure A3. The distribution of 1274 QSOs drawn from our 'good' sample which contain both a Mg II and BOSS pipeline redshift. Top and left panels are histograms of the BOSS pipeline redshift, $z_{\rm pipeline}$ and Mg II derived redshift, $z_{\rm Mg\,II}$. The red dashed curve denotes the one-to-one relation, highlighting a slight offset between the two redshift estimates.

OH sky line features. Typically, the Mg II emission line is one of the strongest anchors in which the source redshift can be determined. At $2.25 \lesssim z \lesssim 2.33$, Mg II is observed within a region that has a relatively low sky background with few sky lines, resulting in a higher Mg II S/N. However, at $z \lesssim 2.25$ and $z \gtrsim 2.33$, the Mg II feature appears amongst OH forest regions with very strong sky lines, resulting both in a noisier signal and it being strongly affected by systematics owing to the sky line residuals. Therefore, the weakening of the Mg II information in determining the QSO redshift within the PCA approach at $z \gtrsim 2.33$ is the likely cause of this transition. Due to this bi-modal nature, which appears to be intrinsic only to the $z_{\rm PCA}$ redshift estimator, and since throughout this work we have assumed the distributions of line profiles can all be represented by a Gaussian distribution, we refrain from adopting the PCA redshift within this analysis.

The next three best candidates after $z_{\rm PCA}$ from the Font-Ribera et al. (2013) analysis are $z_{\rm VI}$, $z_{\rm pipeline}$ and $z_{\rm Mg\,II}$[17]. Of these three, $z_{\rm Mg\,II}$ has the smallest bias, followed by $z_{\rm pipeline}$ and $z_{\rm VI}$. In contrast, $z_{\rm VI}$ has the smallest dispersion, followed by $z_{\rm pipeline}$ and $z_{\rm Mg\,II}$. In this work, we do not consider using the visually inspected redshift as: (i) it has the strongest systematic bias of these choices (ii) it is the least objective measure (not entirely automated), requiring intervention to improve the redshift.

The QSOs within our 'good' sample are selected within $2.08 \leq z \leq 2.5$. By adopting, $z_{\rm Mg\,II}$ instead of $z_{\rm pipeline}$, we lose all QSOs above $z > 2.4$, where no $z_{\rm Mg\,II}$ redshift is reported for any of the sources in DR12Q. Ultimately, the loss of sources reduces the total number of QSOs within our sample by ~ 25 per cent. Despite the reduction in sample size, the $z_{\rm Mg\,II}$ and $z_{\rm pipeline}$ samples return almost identical covariance matrices, signifying consistency in both redshift estimates. The major difference between the two samples is the notably stronger correlation in the velocity offsets for the emission line parameters within the $z_{\rm Mg\,II}$. However, this is to be expected owing to the differences in how the redshifts are estimated. For $z_{\rm Mg\,II}$, a single line profile is used to estimate the redshift, whereas in $z_{\rm pipeline}$ several line profiles are compared to a training set to estimate the best candidate redshift. By using several line profiles to determine the source redshift, intrinsic scatter within the source spectrum can cause the redshift estimates to deviate compared to those from a single line measurement such as $z_{\rm Mg\,II}$. Despite $z_{\rm Mg\,II}$ providing the stronger line profile components for the velocity offsets than the $z_{\rm pipeline}$, we refrain from using the Mg II redshifts as these correlations may be artificially high owing to their estimation from only a single line profile. We deem the correlations from $z_{\rm pipeline}$ sample to be more conservative, while also benefitting from the addition of maintaining a larger statistical sample.

Finally, $z_{\rm Mg\,II}$ and $z_{\rm pipeline}$ return relatively similar dispersions, with $z_{\rm Mg\,II}$ being less biased. With the vast majority of the QSOs in our sample having both redshifts, in principle we can attempt to calibrate our $z_{\rm pipeline}$ sample by the $z_{\rm Mg\,II}$ sample. In Figure A3 we provide the $z_{\rm Mg\,II}$–$z_{\rm pipeline}$ parameter space, with the red dashed curve denoting the one-to-one relation, from which one can see a small offset. We recalibrate our sample for this offset by creating a pipeline corrected redshift,

$$z_{\rm pipeline\_corr} = z_{\rm pipeline} + \langle z_{\rm Mg\,II} - z_{\rm pipeline} \rangle, \quad \text{(A1)}$$

for which we recover a small offset of $\sim 2.77 \times 10^{-3}$ (~ 361 km/s at $z \sim 2.3$).

Employing this BOSS pipeline corrected redshift, we

[17] We do not consider either of $z_{\rm C\,IV}$ or $z_{\rm C\,III]}$ as these have the largest dispersion and bias.

perform the full analysis as performed within Sections 4 and 5. With respect to the covariance matrices, we recover effectively identical correlations amongst our line profile parameters. In order to better gauge the performance of these two possible samples, we recover the fraction of QSOs with a reconstructed flux recovered to within 15 per cent at both 1205Å and 1220Å (see e.g. Figure 11). Though the differences are minor, we find that the original BOSS pipeline redshift ($z_{\rm pipeline}$) performs marginally better than $z_{\rm pipeline\_corr}$. At 1205Å, we recover the QSO flux to within 85 (82) per cent for $z_{\rm pipeline}$ ($z_{\rm pipeline\_corr}$) and 90 (88) per cent at 1220Å[18]. Given that the primary aim of this work is to reconstruct Lyα emission profiles, then the most natural choice for a redshift estimate is the one which performs best at this reconstruction; the line parameter correlations are a secondary product. Therefore, we conclude that owing to the marginally improved statistics on the recovery of the QSO flux from the $z_{\rm pipeline}$ sample relative to $z_{\rm pipeline\_corr}$, and the reasons mentioned previously for not considering other redshift estimates, the best redshift to adopt within this work is $z_{\rm pipeline}$.

APPENDIX B: PREFERENCE FOR THE NUMBER OF GAUSSIAN COMPONENTS FOR A LINE PROFILE

In the construction of our covariance matrix in Section 4.1, we model the Lyα and C IV emission lines as a double component Gaussian whereas for the Si IV +O IV] and C III] lines we consider only a single component Gaussian. In this appendix we detail quantitatively the reasoning for our choices. To objectively differentiate between the various combinations we use the Bayes Information Criterion (BIC; Schwarz 1978; Liddle 2004). The BIC can be computed using,

$$\mathrm{BIC} = -2\ln(\mathcal{L}) + k\ln(N), \tag{B1}$$

where $\mathcal{L}$ is recovered from the χ^2 fit to the individual emission line profile, k is the number of free parameters and N is the number of available data points. A lower recovered BIC for an individual fit to the emission line profile is considered to be a better characterisation of the line.

The comparison between two possible sets of parameters to describe the emission line profile can be quantitatively described by the difference in the two BICs, ΔBIC. A recovered $\Delta\mathrm{BIC} > 10$, provides very strong evidence for the model with the lower BIC, whereas for $6 < \Delta\mathrm{BIC} < 10$ strong evidence exists. Throughout this appendix we discuss each emission line profile individually, and recover the ΔBIC for each model to distinguish between our preferred choices.

B1 The Lyα profile

For the Lyα emission line profile, we consider the full emission line complex within the rest-frame wavelength range $\lambda = [1170, 1280]$ Å. This provides three different potential line species to be fit, Lyα, N V and the Si II line. Given the importance of the N V line for the full line complex, we always fit the N V line in addition to the Lyα line. Therefore, we consider four different combinations to potentially characterise the full Lyα line complex. Firstly, we consider the Lyα line as a double component Gaussian, with a single component Gaussian each for the N V and Si II emission lines. Next, the same, but only a single component Lyα component. Finally, we consider these same two combinations excluding the Si II line.

In Table B1, we report the ΔBIC for these four combinations of emission lines. We compute the BIC as outlined in Equation B1, recovering the χ^2 within the wavelength range $\lambda = [1170, 1280]$ Å. In the construction of Table B1, we use the 'conservative' sample consisting of 2653 QSOs. We then recover the ΔBIC from the difference between a model with a double component Gaussian for Lyα and the single components for N V and Si II. It is clear from this table that this is the preferred choice for modelling the Lyα line complex. Importantly, we find strong evidence for a double component Lyα emission line in over 80 per cent of the QSOs we consider. In principle, we could have explored whether further components for the Lyα line are preferred, however, for the sake of simplicity we have refrained from doing so. Furthermore, while the Si II emission line is typically only a weak ionisation line, and in some cases not resolved within individual QSOs, we still find strong evidence for its inclusion in the QSO fitting (preferred in ~ 62 per cent of our sample).

B2 The Si IV +O IV] profile

Unlike the Lyα line considered above, the blended Si IV +O IV] doublet is a strong, isolated emission line. We therefore consider the simplified choice of whether a single or double component Gaussian is preferred for this line within the wavelength range $\lambda = [1360, 1440]$ Å. From Table B1, we find that a single component Gaussian is very strongly preferred at ~ 65 per cent, which becomes ~ 72 per cent when the criteria is lowered to only strong evidence.

B3 The C IV profile

For the C IV doublet, we consider the emission line complex within the wavelength range, $\lambda = [1500, 1600]$ Å. Again, the C IV emission line is an isolated, very strong high ionisation line, therefore we only need to consider whether a single or double component Gaussian is preferred. In Table B1, we recover very strong evidence for the choice of a double component Gaussian, with this being preferred 94 per cent of the time.

B4 The C III] profile

The C III] emission line is not an isolated line, instead it is the strongest line in a complicated emission line profile (e.g. Vanden Berk et al. 2001). Surrounding the C III] line are contributions from the Fe pseudo-continuum, as well as contributions from the Al III and Si III] line. In this work we consider the contribution of the Fe pseudo-continuum to be negligible to the C III] line, and for all intents and purposes we additionally consider the Si III] contribution negligible.

[18] For completeness, for $z_{\rm Mg\,II}$ we recover the QSO flux to within 15 per cent for 80 per cent of the QSOs at 1205Å and 85 per cent at 1220Å.

Table B1. We use the Bayes Information Criterion (BIC) to distinguish whether the emission lines used within our covariance matrix should be fit with a single or a double component Gaussian. For each potential combination we present the ΔBIC to distinguish between very strong evidence (ΔBIC > 10) and strong evidence (6 < ΔBIC < 10) for a certain combination of lines (first column) relative to an alternative combination (second) column. See text for further details regarding the various combinations of line profiles considered.

Line fitting type (preferred option)	Line fitting type (compared against)	QSOs with ΔBIC > 10 (per cent)	QSOs with 6 < ΔBIC < 10 (per cent)
Lyα double component, N v and Si ii	Lyα double component and N v	60.9	1.1
Lyα double component, N v and Si ii	Lyα single component, N v and Si ii	79.4	0.7
Lyα double component, N v and Si ii	Lyα single component and N v	93.0	0.4
Si iv +O iv] single component	Si iv +O iv] double component	64.7	6.9
C iv double component	C iv single component	94.0	0.3
C iii] single component and Al iii	Only C iii] double component	53.8	2.9
C iii] single component and Al iii	C iii] double component and Al iii	45.7	6.3
C iii] single component and Al iii	Only C iii] single component	92.2	0.9

This latter point arises from the insufficient S/N of these BOSS QSOs to fully resolve and separate out this line contribution. While it may be present in the highest S/N QSOs in our sample, for the vast majority it cannot be easily distinguished. Furthermore, from Vanden Berk et al. (2001), the relative line fluxes differ by over a factor of a hundred.

We consider four scenarios to characterise the C iii] line complex. A double component C iii] line in combination with a single component Gaussian for Al iii. A single component C iii] and Al iii line, and then just a single or double component Gaussian for the C iii] line, without considering the Al iii line. In Table B1, we find just a single C iii] emission line to be very strongly disfavoured with respect to a single component C iii] and Al iii (∼ 92 per cent). However, in the case of the remaining scenarios it is not as clear.

Comparing the single C iii] and Al iii components to a purely double component C iii] (no Al iii), we find marginally stronger evidence for the single C iii] scenario (preferred ∼ 54 per cent of the time). However, note that in allowing for a double component C iii], these two scenarios are effectively the same. For example, the broad component can just fit the Al iii line, leaving a single C iii] line. This degeneracy therefore tends to draw these two scenarios to be more equally preferred. As a result, the single C iii] and Al iii scenario is considered to be preferable.

In the case of comparing a single or double component C iii] line in combination with Al iii, we find comparable results for the ΔBIC. Between these two scenarios we find a single component very strongly preferred in ∼ 46 per cent of the QSOs sampled, which increases to ∼ 52 per cent if we lower the criteria to just strong evidence. It is therefore, not immediately clear from the ΔBIC which is the preferred fitting scenario to consider for the C iii] line complex. As a result, we choose to consider a single C iii] component with a Al iii contribution purely on the grounds of there being less model parameters, removing potential degeneracies between a broad and narrow component C iii] from our covariance matrix.

APPENDIX C: QSO QUALITY ASSESSMENT EXAMPLES

In Section 3.5, we briefly outlined the visual quality assessment we performed to obtain the two QSO samples for the covariance matrix. Within this appendix we provide a more detailed explanation of the respective criteria we consider, as well as providing a select few examples which best indicate the various decisions which we have made.

C1 Removal of QSOs

The first step of our quality assessment was the removal of any QSOs from which we recovered a poor characterisation of the QSO continuum. In Figure C1 we provide two examples of QSOs which were removed failing this criteria. In the top panel, the recovered QSO returns a positive spectral index, α_λ, which is counter to the negative spectral indices recovered by the vast majority of our QSO sample. The recovery of a positive spectral index is unphysical, and likely arises from flux calibration errors. Only a few QSOs within our sample are found to have a positive spectral index. In the bottom panel of Figure C1, we provide an example of a QSO where a single power-law is insufficient to characterise the QSO continuum within our fitting region. To adequately characterise this QSO, a double component continuum pivoting on C iv would be required. The rarity in recovering such an object likely points to this being an error in the data reduction pipeline, or the flux calibration.

The next step in our visual QSO quality assessment is the removal of any QSO spectra which have either (i) missing sections of the observed QSO flux which overlap with our chosen emission lines (ii) a large number of absorption features (iii) either an intervening dense neutral absorber or sufficiently strong/broad absorption blueward of C iv (iv) a Lyα line region that is not well fit or characterised by our MCMC pipeline. In Figure C2 we present nine examples of QSOs removed from our sample. We discuss each one of these QSOs below[19]:

[19] Note that in Figures C2 – C4 we represent the corresponding fitted absorption profiles as being subtracted from the continuum flux level rather than the true flux level. As a result, in some cases

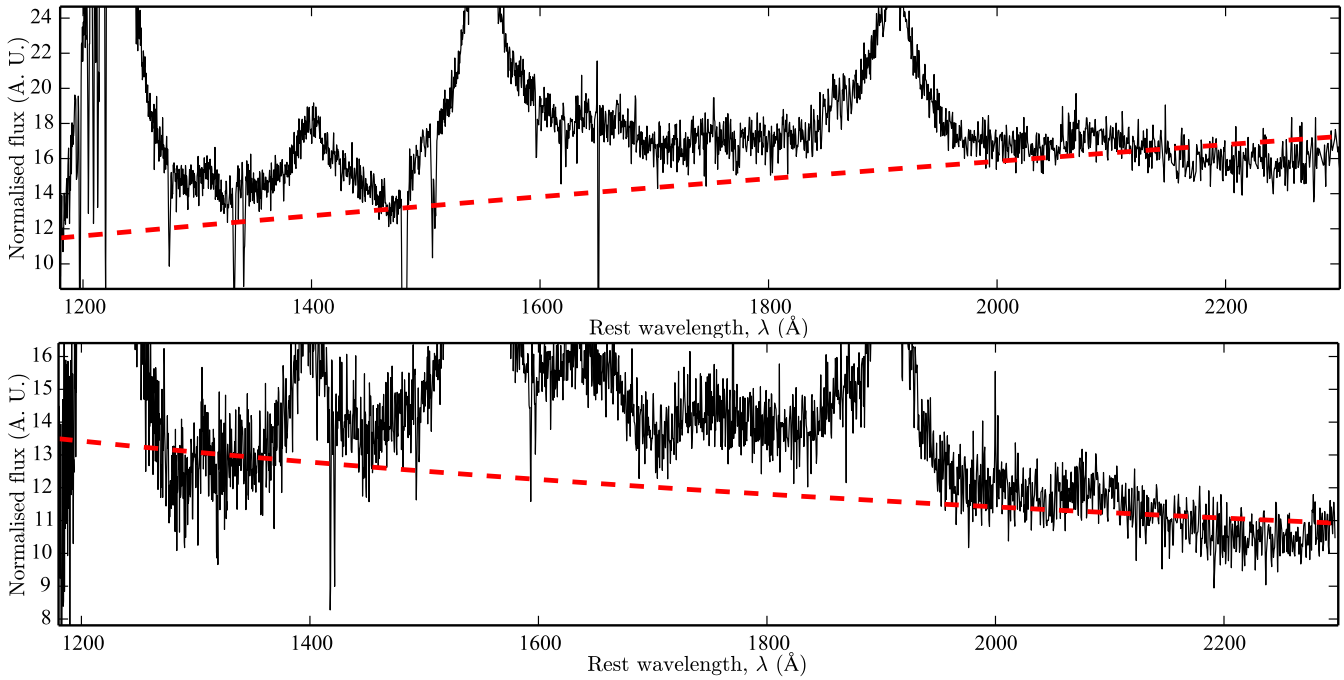

Figure C1. Representative examples of QSOs removed from our samples due to continuum estimation issues. *Top panel:* The recovery of a positive QSO continuum spectral index ($f_\lambda \propto \lambda^{\alpha_\lambda}$). A small number of QSOs contain a positive spectral index, counter to the vast majority of QSOs which have a negative spectral index. *Bottom panel:* A QSO continuum where a single power-law continuum cannot fit the observed spectrum. In this example, a two component continuum with a break near C IV would improve the fit. Only a rare few QSOs show this behaviour, therefore there is no motivation to retain them. Note, 1 A. U. $= 10^{-17}\,\mathrm{erg\,cm^{-2}\,s^{-1}\,\AA^{-1}}$.

a) A clear, dense neutral absorber in the direct vicinity of the QSO. Presence of the neutral absorber completely removes any ability to recover a Lyα broad component.

b) Two strong absorption features near Lyα. The first, blue ward of Lyα removes almost half of the Lyα line. Even in the presence of this absorber, the intrinsic Lyα line profile recovery is still reasonably good. However, it still under predicts the true profile. Furthermore, absorption on the N V line centre removes any ability to recover the N V line, resulting in a much broader Lyα component than should be recovered.

c) Strong Lyα absorption at the line centre. Without a sufficient fraction of the Lyα line centre, the intrinsic Lyα line cannot be recovered. While a broad Lyα component can still be estimated, no narrow component can be reliably obtained.

d) Another, even more extreme example of strong absorption on Lyα line centre. Presence of several strong Lyα absorption features completely removes any information to obtain the Lyα narrow line component.

e) A Lyα line profile with no prominent narrow line component. As a result, the two Gaussian components both prefer to be broad, and are therefore degenerate. The absorption on line centre likely contaminated the Lyα peak determination leading to the two preferred degenerate broad components.

f) A large, connected series of narrow absorption features centred on Lyα. Since the entire Lyα profile is contaminated by absorption, measuring the Lyα narrow component peak amplitude becomes degenerate with the depth of the absorption features.

g) The MCMC pipeline struggles to recover the extremely narrow, sharp peak around Lyα line centre. Evidence of two small absorption features (which are not identified by our algorithm) on either side of line centre appear to cause this peak to become too narrow to fit.

h) Significant blue ward absorption features near the C IV line contaminates the recovery of the broad component. Not a BAL QSO, however the absorption features are still sufficient to contaminate the recovery of the full line profile.

i) Deep, narrow absorption features on the C IV line profile. The unfortunate location of these features at the wing of the C IV narrow component causes a narrower C IV line than should be expected. The depth of these features causes this behaviour, weaker absorption would not impact the narrow component recovery.

it appears that the absorption extends to negative flux. However, this choice is purely to aid the visual representation of the full emission profile fits.

C2 Examples of QSOs in the 'conservative' sample

In Figure C3 we provide three examples of QSOs which are classified as 'conservative' only (not considered for the 'good' sample). In the first example (left panel), several strong narrow absorption features are detected and fit. These are offset from the line centre of Lyα, allowing the peak height of Lyα to still be reliably estimated. Though the intrinsic profile appears to be well recovered, the uncertainties in the peak amplitude owing to potential degeneracies with the relative strengths of the absorption features and the number of iden-

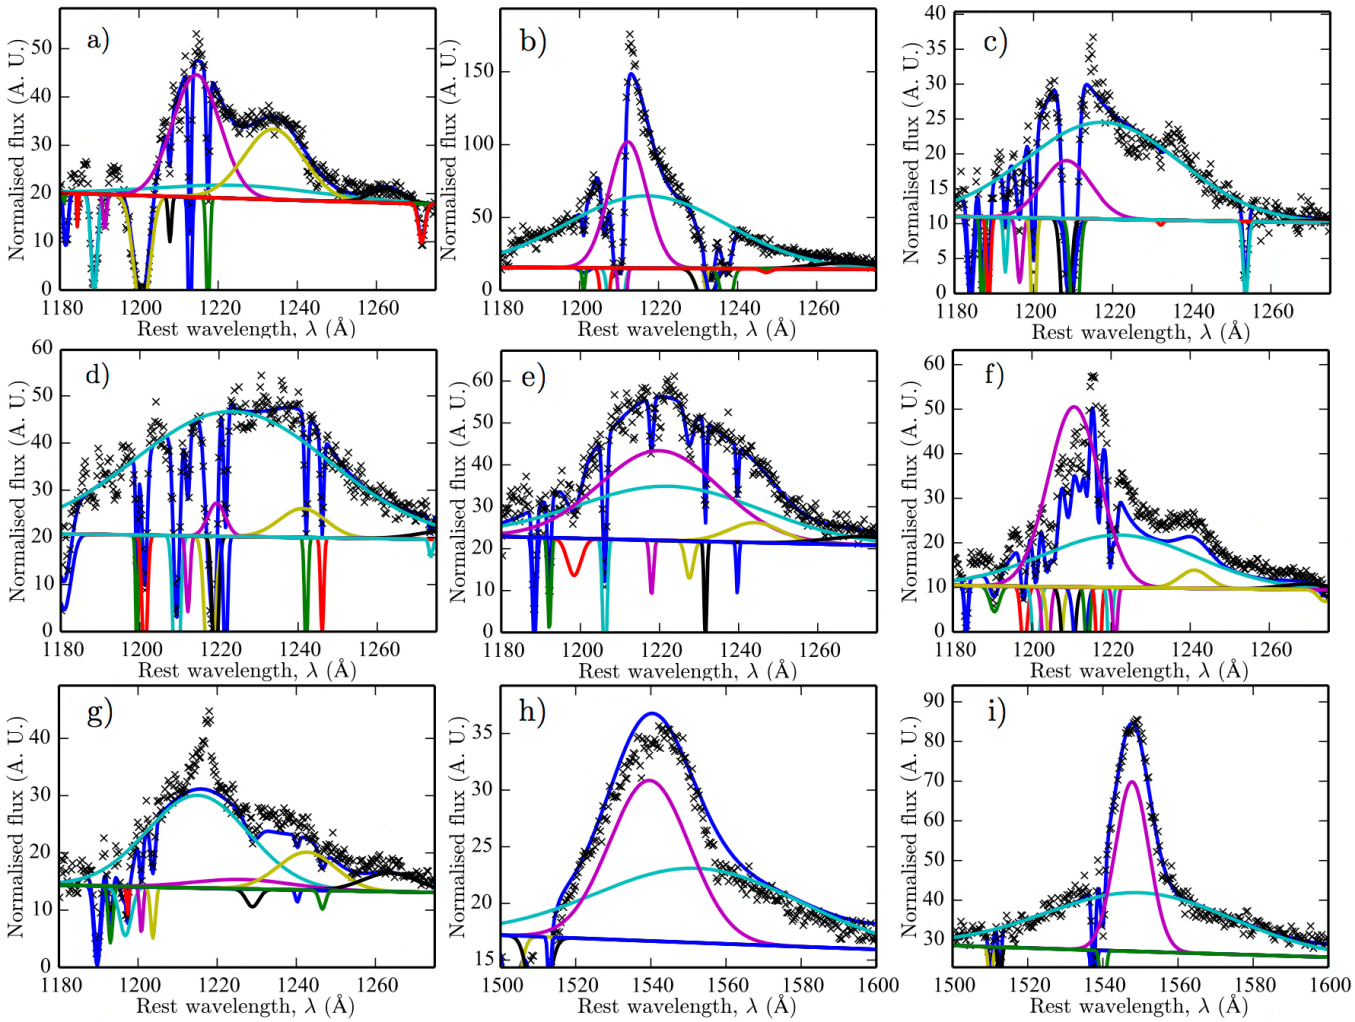

Figure C2. Representative examples of QSOs removed from our samples due to a variety of issues. *Top left:* Intervening strong Lyα absorber. *Top centre:* Strong Lyα absorption offset from line centre impacting the intrinsic profile recovery. *Top right:* Strong Lyα absorption at line centre. *Middle left:* A substantial number of Lyα absorption features contaminating the Lyα emission line profile. *Middle centre:* No prominent Lyα emission profile. *Middle right:* Significant Lyα line absorption which becomes degenerate with the Lyα emission line profile. *Bottom left:* The MCMC algorithm is unable to fit the intrinsic Lyα profile owing to absorption impacting the Lyα emission line profile. *Bottom centre:* C IV absorption in the line wing, impacting the broad component determination. *Bottom right:* Small C IV absorption feature which artificially prefers a narrower line component than should be recovered.

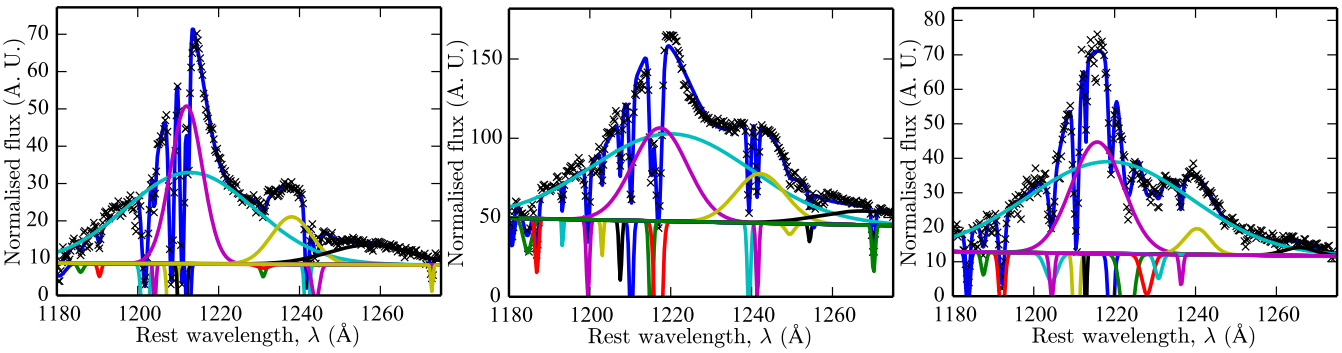

Figure C3. Representative examples of Lyα profiles which are classified in the 'conservative' sample (more uncertainty in the recovery of the intrinsic Lyα profile). *Left panel:* A large number of clustered absorption lines blue-ward of Lyα which are identified and well fit. The pristine red-side of the Lyα line should enable the intrinsic Lyα profile to be well characterised. *Centre panel:* A large number of absorption lines, which are well fit and characterised. The broader Lyα absorption feature will likely impact the recovery of the intrinsic profile, but there is still sufficient evidence that the intrinsic Lyα profile might be well estimated. *Right panel:* A large number of absorption features on both the blue and red side of Lyα. The absence of line centre absorption should enable the intrinsic profile to be recovered with some small uncertainty.

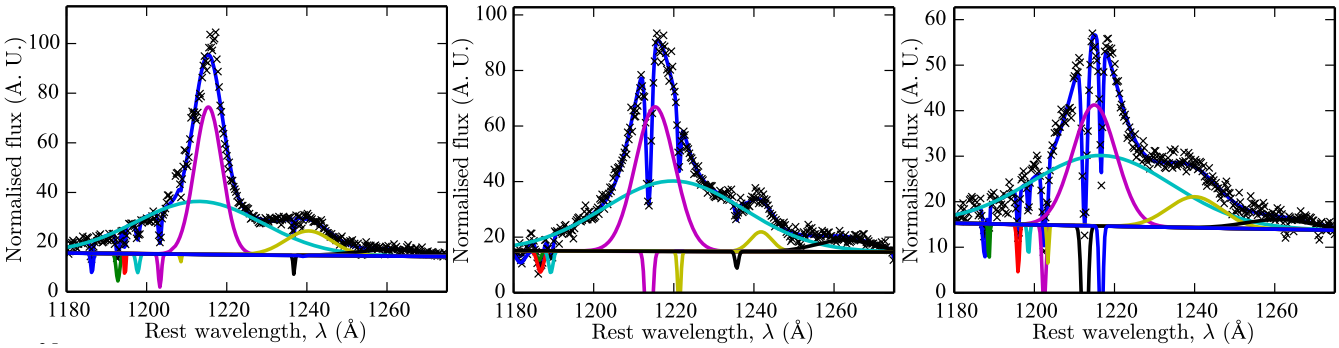

Figure C4. Representative examples of Lyα profiles which are classified in the 'good' sample. *Left panel:* The absence of any absorption. *Centre panel:* small absorption, offset from the centre of Lyα. These features are identified and fit, and therefore do not impact on the ability to fit the intrinsic Lyα profile. *Right panel:* A larger number of absorption feature, all identified and well fit. The narrow width of the absorption features around Lyα should not impact the fitting of the Lyα profile.

tified absorption features causes this to be only characterised as 'conservative'.

The second example (middle panel) has a strong absorption feature on Lyα line centre. However, the relative strength of this feature is not sufficient to completely contaminate the peak of the Lyα line (unlike similar examples in Figure C2). It appears the full Lyα line profile is recovered well, despite numerous narrow absorption features and the strong central absorption. The uncertainties surrounding the narrow Lyα component determination from the strong absorption results in its assessment as 'conservative' only. Finally, in the last example (right panel), the full Lyα line profile is contaminated by a large number of narrow absorption features. None of these are located at line centre, nor do they appear to impact the recovery of the Lyα line profile. However, the sheer number of features causes this QSO to be classified as conservative.

C3 Examples of QSOs in the 'good' sample

Finally, in Figure C4 we present three examples of QSOs explicitly selected for our refined 'good' sample (all QSOs in the 'good' sample appear in the 'conservative' sample). In the first example (left panel), there are only a few very small absorption features. In the absence of anything stronger, these QSOs are classified as 'good'. The second example (middle panel), once again has very few absorption features. Despite being strong, narrow absorption lines, these are well characterised and do not impact the recovery of the intrinsic Lyα line at all. In the final example (right panel), a larger number of absorption lines are present. Importantly though, these are located offset from the Lyα line centre, and therefore we are confident that the true profile is recovered. Note, this profile is similar to the first example in Figure C3. The reason this QSO is classified as 'good' and the previous was 'conservative' arises from the number of absorption lines. In this example, only two narrow features impact the Lyα line, whereas in the previous case, several lines have resulted in a smaller fraction of the total Lyα profile being unaffected by the absorption features. Ultimately, this is only a minor difference that does not change any of the conclusions of this work, since in Section 4.4 we found similar correlation matrices between the two samples.

APPENDIX D: RECOVERY OF CONTINUUM PARAMETERS FROM $\lambda > 1275$Å

In Section 5.2, we alluded to an assumption regarding the usage of the QSO continuum between the full MCMC fitting and the reconstruction profile fitting. Specifically, we assume that the continuum measured from a Lyα obscured or high-z QSO in the region ($\lambda > 1275$Å) will be equivalent to the continuum recovered from fitting the full QSO within the range 1180Å $< \lambda <$ 2300Å. However, by not fitting in the 1180Å $< \lambda <$ 1275Å region, we could be losing information from the Lyα line profile (and surrounding region) that could be otherwise used to refine the estimate of the QSO continuum. Of course, this cannot be tested on an actual QSO that requires reconstruction, but instead we can statistically test this assumption across the full QSO sample we have used within this work.

In Figure D1, we present the 2D scatter of the two QSO continuum parameters, the continuum flux amplitude normalised to 1450 Å (f_{1450}, left panel) and the spectral index, α_λ (right panel). To facilitate the comparison, the red dashed curve corresponds to the one-to-one relation. Note that while the flux amplitude is normalised at 1450Å, which should be unaffected by the different fitting regions, we allow a perturbation on this flux to refine its true normalisation. In the case of this continuum normalisation, we find very little scatter, and all QSOs lie along the one-to-one relation indicative of their being no adverse effects from this assumption.

In the case of the QSO spectral index, we again find all QSOs to essentially lie perfectly along the one-to-one relation. However, in several instances, there are QSOs which deviate away from this line. This small scatter indicates that for a select few QSOs, fitting in the 1180Å $< \lambda <$ 1275Å region is important for the recovery of the true continuum spectral index. Importantly, around Lyα even for the extreme QSOs in this 2D plot, the level of discrepancy in the continuum flux will be at most 10 per cent. Furthermore, this is purely for the continuum flux, and not the total line flux. Referring to Figure 9, a 10 per cent error on the continuum flux at Lyα will fall well within the errors of the reconstructed Lyα emission line profile. Therefore, within the model uncertainties of the Lyα reconstruction pipeline, this scatter is not significant to drastically affect the recon-

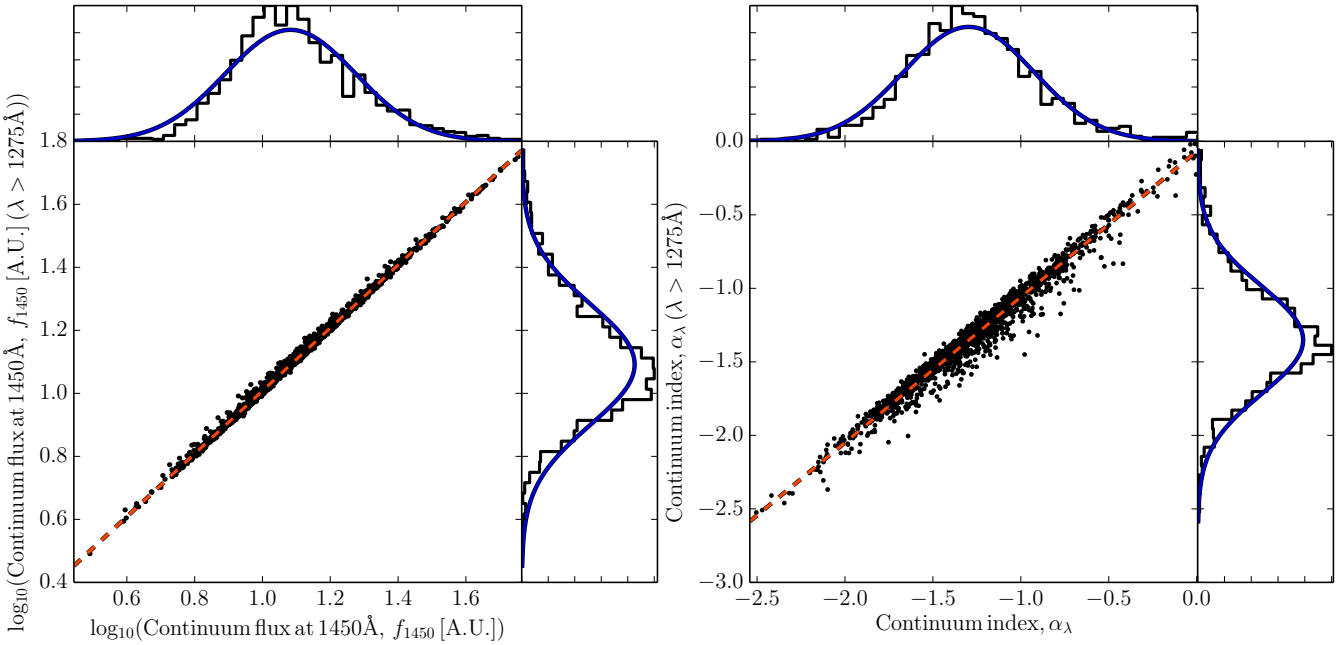
log10(Continuum flux at 1450Å, f1450 [A.U.]) (λ > 1275Å))
log10(Continuum flux at 1450Å, f1450 [A.U.])
Continuum index, αλ (λ > 1275Å)
Continuum index, αλ

Figure D1. 2D scatter plots comparing the estimation of the two continuum parameters between the full MCMC fit and the Lyα masked QSO fit for reconstruction (masked at $\lambda < 1275$Å). *Left panel:* The continuum flux normalised at 1450Å, f_{1450}. *Right panel:* The continuum spectral index, α_λ. The red dashed curve corresponds to the one-to-one relation. There is very little scatter about the one-to-one relation, indicating there are no biases in the estimation of the continuum parameters when fitting either the full QSO, or fitting at $\lambda > 1275$Å. Note, 1 A. U. $= 10^{-17}\,\mathrm{erg\,cm^{-2}\,s^{-1}\,Å^{-1}}$.

structed Lyα line profile and hence the assumption on using the same QSO continuum is justified.